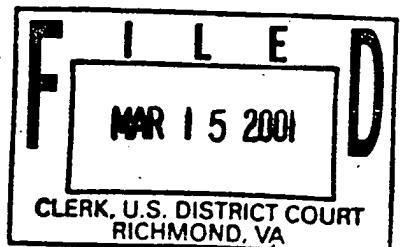


IN THE UNITED STATES DISTRICT COURT
FOR THE EASTERN DISTRICT OF VIRGINIA
Richmond Division



RAMBUS, INC.,

Plaintiff,

v.

Civil Action No. 3:00cv524

INFINEON TECHNOLOGIES AG
and INFINEON TECHNOLOGIES
NORTH AMERICA CORP.,

Defendants.

MEMORANDUM OPINION

This action involves four patents and fifty-seven claims. All four patents in suit descend from a common progenitor, the specification of which controls the patents in suit. The parties are in agreement that construction of the claims here at issue is confined to construction of eight disputed terms ("bus," "block size," "read request," "write request," "transaction request," "first external clock signal," "second external clock signal" and "integrated circuit device") each of which, with but one exception,¹ has the same meaning in each claim in issue in all four patents in suit. Hence, the agreed upon scope of claim construction is to construe the eight terms.

¹ The parties agree that all the terms have the same meaning throughout with the exception of "integrated circuit device." The Defendants contend that this term has a different meaning in one patent due to representations made to the Patent and Trademark Office during the prosecution of that patent.

The parties have briefed the issues, have presented evidence at a hearing conducted pursuant to the requirements of Markman v. Westview Instruments, Inc., 517 U.S. 370 (1996), and have argued orally. Against this background, the eight disputed terms, and hence the claims, are accorded the constructions set forth below.

BACKGROUND

In 1990, the co-founders of Rambus, Inc. ("Rambus"), Mark Horowitz and Paul Michael Farmwald, filed a patent application describing numerous inventions designed to increase the operating speed of memory devices in computers. The Patent Office determined that this application, U.S. Patent App. No. 07/510,898 ("the '898 application"), actually contained 11 independent and distinct inventions, required Rambus to select only one of those inventions to pursue in the '898 application, and allowed Rambus to file divisional applications on the remaining inventions described in the '898 application. Rambus did precisely that, electing to pursue one invention within the '898 application and thereafter filing ten more applications in the next six months. Subsequently, continuation and divisional applications were filed on these ten applications; and thus, to date, Rambus has been granted 31 patents based on the 1990 '898 application. Numerous applications are currently pending.

By way of background, the patented inventions have to do with computer memory devices called Dynamic Random Access Memory

("DRAM"), and a system and devices for increasing the speed at which data or information is transferred between the DRAM and the Central Processing Unit ("CPU") of a computer. The DRAM is a high-speed, short-term memory device where information being used by the CPU is stored. The patents in suit describe numerous inventions respecting the memory interface and a new type of "bus" which carries information or data. The "Field of Invention" section of the specification, common to all patents in suit, gives the following overview of the inventions:

[a]n integrated circuit bus interface for computer and video systems is described which allows high speed transfer of blocks of data, particularly to and from memory devices, with reduced power consumption and increased system reliability. A new method of physically implementing the bus architecture is also described.

U.S. Patent No. 6,034,918 (issued March 7, 2000) ("the '918 patent"), col. 1, ll. 20-25.²

On August 8, 2000, Rambus instituted this action for the infringement of four of its patents against Infineon Technologies AG (a German corporation), Infineon Technologies, Inc. (a German corporation) Infineon Technologies North America Corp. (a Delaware corporation) and Infineon Technologies Holding North America, Corp. (a Delaware corporation) (collectively referred to as "Infineon").

² All the patents in suit, and all the patents springing from the 1990 '898 application, contain the same specification. For ease of citation, all references to the specification will be to the '918 patent.

The first of the patents in suit, U.S. Patent No. 5,953,263 (issued Sept. 14, 1999) ("the '263 patent"), claims a latency invention which involves the use of a programmable register on the DRAM chip to store a value representative of a time delay. The latency invention makes the DRAM response time more predictable because the CPU knows precisely when it will receive data from the DRAM in response to a transaction request, thereby allowing the system to plan for transfers and improving overall traffic flow over the bus.³ Claims 1-5, 14, 16-19, 21, 23-25, 27-28, 30 and 32-33 of the '263 patent are at issue in this action.

Secondly, in U.S. Patent No. 5,954,804 (issued Sept. 21, 1999) ("the '804 patent"), Rambus claims a delayed lock loop (DLL) on a DRAM chip, which allows precise timing of the output of data. In essence, the DLL allows the DRAM chip to collect the data from the memory cells and then paces the release of that information over the bus. The DLL becomes useful when operating the DRAM at high

³ Claim 1 of the '263 patent is representative of this invention:

1. A synchronous semiconductor memory device having at least one memory section which includes a plurality of memory cells, the memory device comprises:

a programmable register to store a value which is representative of a delay time after which the memory device responds to a read request.

rates of speed.⁴ Claim 26 of the '804 patent is the only claim involving this invention at issue in this action.

The third patent, U.S. Patent No. 6,034,918 (issued Mar. 7, 2000) ("the '918 patent"), covers the variable block size invention, which involves the use of circuitry to allow for the output of variable-sized blocks of data over the bus in response to a transaction request. The additional circuitry allows a user, such as a CPU, to select differing sizes or blocks of data, instead

⁴ Claim 26 of the '804 patent describes DLL in combination with the latency invention:

26. An integrated circuit device having at least one memory section which includes a plurality of memory cells, wherein the integrated circuit device outputs data on an external bus synchronously with respect to first and second external clock signals, the integrated circuit device comprises:

a first internal register to store a value which is representative of a number of clock cycles to transpire before the integrated circuit device responds to a read request;

delay locked loop circuitry to generate an internal clock signal using the first and second external clock signals; and

interface circuitry, coupled to the external bus to receive a read request, the interface circuitry includes a plurality of output drivers, coupled to the external bus, to output data on the external bus in response to the internal clock signal, synchronously with respect to the first and second external clock signals and in accordance with the value stored in the first internal register.

of a single piece of data.⁵ Claims 1-2, 6, 8-9, 13, 15-20, 24-25, 29-31, 33 and 34 of the '918 patent are at issue in this action.

Lastly, U.S. Patent No. 6,032,214 (issued Feb. 29, 2000) ("the '214 patent") claims double data rate ("DDR") as the invention. In general, memory devices send and receive information according to a clock contained within the computer system. Clocks are a common, but important, feature of all computer systems. Before the DDR invention, information was transferred only on the "tick" of the clock. The memory device using that type of transfer regulator is called a Synchronous DRAM, or "SDRAM." The DDR invention allows information from the SDRAM to be sent out on both the "tick" and

⁵ Claim 18 of the '918 patent describes this invention as:

18. A method of operation of a synchronous memory device, wherein the memory device includes a plurality of memory cells, the method of operation of the memory device comprises:

receiving an external clock signal;

receiving first block size information from a bus controller, wherein the first block size information defines a first amount of data to be output by the memory device onto a bus in response to a read request;

receiving a first request from the bus controller; and

outputting the first amount of data corresponding to the first block size information, in response to the first read request, onto the bus synchronously with respect to the external clock signal.

the "tock" (or the rising and falling edges) of the computer's internal clock, thereby doubling the data output of the SDRAM for a given clock rate.⁶ Claims 1-2, 4, 6, 9-11, 14-16, 18-19, 21, 24-26 and 29 of the '214 patent are at issue in this action.

Infineon makes, uses, sells or offers to sell, and imports SDRAM devices, DDR SDRAM devices and Synchronous Graphics RAM ("SGRAM") devices, as well as products, such as computers, servers, automated teller machines, telephones and telephone systems and point of sale terminals, all of which contain SDRAM, DDR SDRAM or

⁶ Claim 15 of the '214, which is representative of this invention, covers this invention in combination with the variable block size described in the '918 patent:

A method of operation of a synchronous memory device, wherein the memory device includes a plurality of memory cells, the method comprising:

receiving first block size information, wherein the first block size information defines a first amount of data to be output onto a bus in response to a read request;

receiving a first read request; and

outputting the first amount of data corresponding to the first block size information, in response to the first read request, onto the bus synchronously with respect to a first and a second external clock signal wherein a first portion of the first amount of data is output synchronously with respect to the first external clock signal and a second portion of the first amount of data is output synchronously with respect to the second external clock signal.

SGRAM devices. Rambus alleges that all of those devices and the products and modules into which they are incorporated infringe some or all of the patents in suit. Infineon denies that its products infringe any of those patents.

DISCUSSION

I. The Legal Standard

Patent infringement analysis involves two steps: ascertaining the proper construction of the patent claim and determining whether the accused method or product infringes the properly construed claim. Vitronics Corp. v. Conceptronic, Inc., 90 F.3d 1576, 1582 (Fed. Cir. 1996). A patent contains two distinct elements: "First, it contains a specification describing the invention 'in such full, clear, concise and exact terms as to enable any person skilled in the art . . . to make and use the same.' 35 U.S.C. § 112 Second, a patent includes one or more 'claims,' which 'particularly point[] out and distinctly claim the subject matter which the applicant regards as his invention.' Markman v. Westview Instr., Inc., 517 U.S. 370, 373 (1996).

The construction or interpretation of a claim is a question of law. Markman v. Westview, 52 F.3d 967, 979 (Fed. Cir. 1995) (en banc), aff'd 517 U.S. 370 (1996). "[I]n interpreting an asserted claim, the court should look first to the intrinsic evidence of record, i.e., the patent itself, including the claims, the

specification and, if in evidence, the prosecution history. Such intrinsic evidence is the most significant source of the legally operative meaning of disputed claim language." Vitronics, 90 F.3d at 1582 (internal citations omitted). If the intrinsic evidence is insufficient to resolve ambiguity in the meaning of claims, the court may rely upon extrinsic evidence to understand the technology and to construe the claims. Id. at 1584. "Extrinsic evidence is that evidence which is external to the patent and file history, such as expert testimony, inventor testimony, dictionaries, and technical treatises and articles." Id. Extrinsic evidence, however, may not be used to contradict the claim language or the meanings established in the specification. Id. "Any other rule would be unfair to competitors who must be able to rely on the patent documents themselves, without consideration of expert opinion that then does not even exist, in ascertaining the scope of a patentee's right to exclude." Id. (quoting Southwall Tech. Inc. v. Cardinal IG Co., 54 F.3d 1570, 1578 (Fed. Cir. 1995), cert. denied, 516 U.S. 987 (1995)).

In the examination of the intrinsic evidence, "there is a hierarchy of analytical tools. The actual words of the claim are the controlling focus." Digital Biometrics, Inc. v. Identix, Inc., 149 F.3d 1335, 1344 (Fed. Cir. 1998). Thus, a court should first "look to the words of the claims themselves, both asserted and nonasserted, to define the scope of the patented invention."

Vitronics, 90 F.3d at 1582. See Pitney Bowes, Inc. v. Hewlett-Packard Co., 182 F.3d 1298, 1305 (Fed. Cir. 1999) ("The starting point for any claim construction must be the claims themselves."); K-2 Corp. v. Salomon S. A., 191 F.3d 1356, 1362 (Fed. Cir. 1999) ("We begin, of course, with the language of the claims").

"The general rule is that terms in the claim are to be given their ordinary and accustomed meaning." Id. See also Vitronics, 90 F.3d at 1582. "It is the person of ordinary skill in the field of the invention through whose eyes the claims are construed. Such person is deemed to read the words used in the patent documents with an understanding of their meaning in the field, and to have knowledge of any special meaning and usage in the field." Multiform Desiccants, Inc v. Medzam, Ltd., 133 F.3d 1473, 1477 (Fed. Cir. 1998). Notwithstanding that terms in the claim and specification are presumed to carry the ordinary meaning that they would have to one of ordinary skill in the field, "a patentee may choose to be his own lexicographer and use terms in a manner other than their ordinary meaning, as long as the special definition of the term is clearly stated in the patent specification or file history." Vitronics, 90 F.3d at 1582. See also Hoescht Celanese Corp. v. BP Chems. Ltd., 78 F.3d 1575, 1578 (Fed. Cir. 1996), cert. denied 519 U.S. 911 (1996) ("A technical term used in a patent document is interpreted as having the meaning that it would be given by persons experienced in the field of the invention, unless

it is apparent from the patent and the prosecution history that the inventor used the term with a different meaning").

That is, the ordinary and accustomed meaning of a disputed claim term is presumed to be the correct one, subject to the following. First, a different meaning clearly and deliberately set forth in the intrinsic materials -- the written description or the prosecution history -- will control. Second, if the ordinary and accustomed meaning of a disputed term would deprive the claim of clarity, then further reference must be made to the intrinsic -- or in some cases, extrinsic -- evidence to ascertain the proper meaning. In either case, a party wishing to alter the meaning of a clear claim term must overcome the presumption that the ordinary and accustomed meaning is the proper one, demonstrating why such an alteration is required.

K-2 Corp., 191 F.3d at 1362-63 (internal citations omitted). See Hoganas AB v. Dresser Indus., Inc., 9 F.3d 948, 951 (Fed. Cir. 1993) ("Although a patentee can be his own lexicographer, as we have repeatedly said, the words of a claim will be given their ordinary meaning, unless it appears that the inventor used them differently." (internal quotations omitted)). Cf. Johnson Worldwide Assoc., Inc. v. Zebco Corp., 175 F.3d 985, 990 (Fed. Cir. 1999) (indicating that the patentee must set "forth an explicit definition for a claim term"). "Thus, second, it is always necessary to review the specification to determine whether the inventor has used any terms in a manner inconsistent with their ordinary meaning." Vitronics, 90 F.3d at 1582 (emphasis added); CVI/Beta Ventures, Inc. v. Tura LP, 112 F.3d 1146, 1153 (Fed. Cir.

1997), cert. denied 522 U.S. 1109 (1998) (same). See also Toro Co. v. White Consolidated Indus., Inc., 199 F.3d 1295, 1299 (Fed. Cir. 1999) ("words of ordinary usage must nonetheless be construed in the context of the patent documents").

The specification acts as a dictionary when it expressly defines terms used in the claims or when it defines terms by implication. . . . The specification contains a written description of the invention which must be clear and complete enough to enable those of ordinary skill in the art to make and use it. Thus, the specification is always highly relevant to the claim construction analysis. Usually, it is dispositive; it is the single best guide to the meaning of a disputed term.

Vitronics, 90 F.3d at 1582 (emphasis added). The ordinary meaning of claim terms is a "heavy presumption" to be overcome. Johnson Worldwide, 175 F.3d at 989.

As the third category of intrinsic evidence, "the court may also consider the prosecution history of the patent, if in evidence. This history contains the complete record of all the proceedings before the Patent and Trademark Office ["PTO"], including any express representations made by the applicant regarding the scope of the claims." Vitronics, 90 F.3d at 1583 (internal citations omitted). "[A]rguments made during prosecution regarding the meaning of a claim term are relevant to the interpretation of that term in every claim of the patent absent some clear indication to the contrary." Southwall Tech., 54 F.3d at 1579. "The prosecution history limits the interpretation of

claim terms so as to exclude any interpretation that was disclaimed during prosecution." Id. at 1576. "Claims cannot be construed in one way to obtain their allowance and in a different way against accused infringers." Id. See Digital Biometrics, 149 F.3d at 1344 ("The prosecution history is relevant because it may contain contemporaneous exchanges between the patent applicant and the PTO about what the claims mean").

When consideration of these three sources resolves the disputes over the asserted claim terms (as it generally should), reliance on extrinsic evidence to construe the claim is improper. Vitronics, 90 F.3d at 1583. This is because the claims, specification and file history comprise the public record of the patentee's claim, and to allow the public record (upon which competitors are entitled to rely when investigating the scope of the patentee's claimed invention), to be altered or changed by extrinsic evidence is to undermine the notice function of the public record. Id.

The preference for intrinsic evidence, however, does not preclude a court from considering or relying upon extrinsic evidence:

Vitronics does not prohibit courts from examining extrinsic evidence, even when the patent document is itself clear. . . . Moreover, Vitronics does not set forth any rules regarding the admissibility of expert testimony into evidence. . . . Rather, Vitronics merely warned courts not to rely on extrinsic evidence in claim construction to

contradict the meaning of claims discernible from thoughtful examination of the claims, the written description, and the prosecution history--the intrinsic evidence.

Pitney Bowes, 182 F.3d at 1308 (emphasis in original). See also Bell & Howell Document Mngmt. Prods. Co. v. Altek Sys., 132 F.3d 701, 706 (Fed. Cir. 1997) ("Use of expert testimony to explain an invention may be useful. But reliance on extrinsic evidence to interpret claims is proper only when the claim language remains genuinely ambiguous after consideration of the intrinsic evidence. . . .").

This is especially the case with respect to technical terms, as opposed to non-technical terms in general usage or terms of art in the claim-drafting art. . . . Indeed, a patent is both a technical and a legal document. While a judge is well-equipped to interpret the legal aspects of the document, he or she must also interpret the technical aspects of the document, and indeed its overall meaning, from the vantage point of one skilled in the art.

Pitney Bowes, 182 F.3d at 1309.

Within the category of extrinsic evidence, some types of evidence are preferred over others: "prior art documents and dictionaries, . . . are more objective and reliable guides [than expert testimony]. Unlike expert testimony, these sources are accessible to the public in advance of litigation. . . . Indeed, opinion testimony on claim construction should be treated with the utmost caution, for it is no better than opinion testimony on the meaning of statutory terms." Id. at 1585.

These fundamental precepts inform and guide the construction of the claims at issue in this action. As mentioned previously, there are 57 different claims being asserted under the four patents in suit and each of those claims are in dispute and therefore must be construed. However, in their claim construction briefs the parties have circumscribed that rather daunting task by identifying eight terms to be interpreted. At the Markman hearing, the parties agreed that (with a previously noted exception) these eight terms have the same meaning in each of the 57 asserted claims. As a result, the claim construction task in this action reduces to construing the eight disputed terms. That task is undertaken *seriatim*.

II. Claim Construction

A. "Bus"

The parties dispute the meaning of "bus" as that term is used throughout the claims of the patents in suit. Rambus argues that "bus" means any "set of signal lines (for example, wires) to which a number of devices are connected, and over which information is transferred between devices." According to Rambus, "the term "bus" is old and very common in the electrical arts" and, in the patents in suit, the term is used in its ordinary and customary sense "as a set of signal lines over which information is transferred."⁷ To

⁷ Plaintiff Rambus, Inc.'s Markman Brief Concerning Claim Construction, p. 13.

support the contention that this is the ordinary and customary construction of the term "bus," as used in its patents, Rambus relies not upon intrinsic evidence but upon the extrinsic evidence of the *IEEE (Institute of Electrical and Electronics Engineers) Standard Dictionary of Electrical and Electronics Terms*, Fourth Ed., IEEE Inc., New York (1988), p. 116, to explain how one skilled in the art would understand the term. The IEEE Dictionary defines a bus as "a set of signal lines used by an interface system, to which a number of devices are connected, over which information is transferred between the devices." Id.⁸

Infineon, on the other hand, contends that "bus" actually has a specialized meaning conferred by the specification of the patents in suit, which describes and explains the bus and its use with the other inventions as the Rambus "multiplexed bus." Before the '898 application was filed in 1990, most buses generally had point-to-point interfaces wherein the CPU would communicate with different memory devices by different and separate lines. Furthermore, within each bus in the prior art, the lines would be dedicated to

⁸ Of course, a court cannot use an inconsistent dictionary definition to contradict the meaning derived from the intrinsic evidence, but such definition may be of some assistance to the court in interpreting technical terms. See Vanguard Prods. Corp. v. Parker Hannifin Corp., 234 F.3d 1370, 1372 (Fed. Cir. 2001) ("Although a dictionary definition may not enlarge the scope of a term when the specification and the prosecution history show that the inventor, or recognized usage in the field of the invention, have given the term a limited or specialized meaning, a dictionary is often useful to aid the court in determining the correct meaning to be ascribed to a term as it was used.")

carrying either data, address, control or device-select information. In the new inventive Rambus bus, a single bus is multiplexed so that the bus lines carry all the address, control, data and device-select information over a single bus. In Infineon's view, the use of the term "bus" throughout the claims is limited to the new inventive bus described in the specification.

1. The Claim Language

The analysis begins by first considering the claim language.⁹ Most of the 57 claims at issue use the term "a bus" or "the bus" or "an external bus." None of the claims, however, expressly define the term "bus," nor do they dispositively support either proposed definition. Rather, the claims generally speak of outputting or inputting data over a bus.

Infineon urges the court to consider the language of claim 26 of the '918 patent as illustrative of its view of the term:

⁹ The term "bus" is used in claims 1, 2, 6, 8, 16, 18, 19, 20, 24, 33, and 14 of the '918 patent, claims 1, 2, 4, 10, 15, 16, 18, and 25 of the '214 patent, claims 2, 14, 27, and 30 of the '263 patent and claim 26 of the '804 patent.

26. An integrated circuit device having at least one memory section which includes a plurality of memory cells, wherein the integrated circuit device outputs data on an external bus synchronously with respect to first and second external clock signals, the integrated circuit device comprises:

interface circuitry, coupled to the external bus to receive a read request, the interface circuitry includes a plurality of output drivers, coupled to the external bus, to output data on the external bus in response to the internal clock signal, synchronously with respect to the first and second external clock signals and in accordance with the value stored in the first internal register.

'918 patent, Claim 26 (emphasis added). Infineon posits that this claim calls for data to be output onto the bus, and a read request to be received on the same bus, thus supporting its conclusion that "bus" means a multiplexed bus.¹⁰ While the language of this single claim somewhat supports Infineon's construction, the specification must be reviewed to determine how the inventors used the term "bus" and whether they intended the term to have a special meaning. See Watts v. XL Sys., Inc., 232 F.3d 877, 882 (Fed. Cir. 2000) ("One purpose for examining the specification is to determine if the patentee has limited the scope of the claims"). "[E]ven if [the claims] were clear on their face, [the court] must consult the

¹⁰ The testimony of Infineon's expert, Mr. Joseph McAlexander also supports this conclusion. See Markman Hearing, Tr. pg. 370 l. 13 to pg. 371, l. 19 (explaining that claim 1 of the '918 patent clearly indicates that a read request and output data are to travel across a single bus).

specification to determine if the patentee redefined any of those terms." Id. at 883.

2. The Specification

A close study of the patent specification reveals that, not only did the inventors act as their own lexicographers in defining the term "bus" to be the new inventive bus, but they also repeatedly explained how their various inventions worked in conjunction with the new bus, which they describe to be a centerpiece of the systems they claim to have invented.

 The specification clearly and unambiguously describes the bus of the invention to be the inventive multiplexed bus. In the "Summary of Invention" the specification states:

The present invention includes a memory subsystem comprising at least two semiconductor devices, including at least one memory device, connected in parallel to a bus where the bus includes a plurality of bus lines for carrying substantially all address, data and control information needed by said memory devices, where the control information includes device-select information and the bus has substantially fewer bus lines than the number of bits in a single address, and the bus carries device-select information without the need for separate device-select lines connected directly to individual devices.

'918 patent, col. 3, ll. 50-60 (emphasis added). And again, later in the same section, the specification states, "In this system of this invention, DRAMs and other devices receive address and control information over the bus and transmit or receive requested data

over the same bus. Each memory device contains only a single bus interface with no other signal pins." '918 patent, col. 4, lines 9-13 (emphasis added). See also '918 patent col. 3, l. 61 through col. 4 l. 1. (the DRAM "is modified to use a wholly bus-based interface rather than the prior art combination of point-of-point and bus-based wiring used with conventional versions of these devices. The new bus includes clock signals, power and multiplexed address, data and control signals").

Throughout the "Detailed Description," the specification repeatedly explains the use of the new multiplexed bus:

The present invention is designed to provide a high speed, multiplexed bus for communication between processing devices and memory devices and to provide devices adapted for use in the bus system.

The bus consists of a relatively small number of lines connected in parallel to each device on the bus. The bus carries substantially all address, data and control information needed by devices for communication with other devices on the bus. In many systems using the present invention, the bus carries almost every signal between every device in the entire system. There is no need for separate device-select lines since device-select information for each device on the bus is carried over the bus. There is no need for separate address and data lines because address and data information can be sent over the same lines.

Virtually all the signals needed by the computer system can be sent over the bus.

'918 patent, col. 5, ll. 29-45 (emphasis added). The inescapable lesson that emerges from comparing the claims of the patents with the inventors' fulsome textual description of the invention is that the inventions include a new bus and new devices that work with the inventive bus, all to the inventor's stated purpose, which is "to provide a high speed multiplexed bus for communication between processing devices and memory devices and to provide devices adapted for use in the bus system." '918 Patent, col. 5, ll. 29-33 (emphasis added).

Additionally, not only does the specification define "bus" to be a multiplexed bus, but it also sets a background for explaining how the inventive multiplexed bus works with various other features of Rambus' inventions. Thus, the explanation of the inventions also supports the conclusion that the term "bus" means the multiplexed bus. For example, every embodiment described in the specification involves the use of a multiplexed bus.¹¹ Not once do

¹¹ See e.g. '918 patent, col. 4, ll. 1-4 ("In a preferred implementation, 8 bus data lines and an AddressValid bus line carry address, data and control information for memory addresses up to 40 bits wide.") (emphasis added); '918 patent, col. 5, ll. 59-64 ("In the preferred implementation, memory devices are provided that have no connections other than the bus connections described herein and CPUs are provided that use the bus of this invention as the principal, if not exclusive, connection to memory and to other devices on the bus.") (emphasis added); '918 patent, col. 8 ll. 17-25 ("The preferred bus architecture of this invention comprises 11 signals: BusData[0:7]; AddrValid; Clk1 and Clk2; plus an input reference level and power and ground lines connected in parallel to each device The bus lines for BusData[0:7] signals form a byte-wide, multiplexed data/address/control bus").

the patents indicate that any of the inventions can, or should be, used with the prior art dedicated bus architecture.

This understanding is confirmed by the testimony of Mr. Joseph McAlexander, Infineon's expert, who explains that the patents "describe several bus architectures. But in every instance when they describe the bus of the invention it is always a multiplexed address, data and control bus." Markman Hearing, Tr. p. 360, l. 25 to p. 361, l. 4. Rambus' expert did not refute this conclusion.

In Toro Co. v. White Consolidated Indus., Inc., 199 F.3d 1295 (Fed. Cir. 1999), the Federal Circuit found it significant that the disputed patent contained only one embodiment of the invention. On the issue of whether a "ring" described in the patent must be attached to the "cover," the court noted that:

The specification and drawings show the restriction ring as 'part of' and permanently attached to the cover. No other structure is illustrated or described. . . .

. . . This is not simply the preferred embodiment; it is the only embodiment. . . .

. . . Nowhere in the specification, including its twenty-one drawings, is the cover shown without the restriction ring attached to it.

Id. at 1301. See also O.I. Corp. v. Tekmar Co. Inc., 115 F.3d 1576, 1581 (Fed. Cir. 1997) (rejecting patentee's argument that the invention could have smooth or cylindrical walls when "[a]ll of the 'passage' structures contemplated by the written description are thus either non-smooth or conical."); General Amer. Transp. Corp.

v. Cryo-Tran, Inc., 93 F.3d 766, 770 (Fed. Cir. 1996), cert. denied 520 U.S. 1155 (1997) (the disputed claim construction was "not just the preferred embodiment of the invention; it is the only one described. Nothing in the claim language, specification, or drawings suggests that any of the [limitations] may be eliminated . . .") (emphasis in original). Likewise, it is significant here that Rambus does not list a single example of how any of the new inventions would work with any type of bus other than a multiplexed bus.¹²

The Federal Circuit's holding in Wang Labs., Inc. v. America Online, Inc., 197 F.3d 1377 (Fed. Cir. 1999) is instructive as well. In Wang, the court considered whether the ordinary and

¹² The failure of the specification to describe any other kind of bus in connection with the invention distinguishes the principal decision upon which Rambus relies, Johnson Worldwide Assoc., Inc. v. Zebco Corp., 173 F.3d 985 (Fed. Cir. 1999). In Johnson Worldwide, the Federal Circuit placed great emphasis upon the fact that the disputed claims did not require the narrower construction. 175 F.3d at 991. The Johnson Worldwide opinion distinguished Laitram Corp. v. Morehouse Indus., Inc., 143 F.3d 1456 (Fed. Cir. 1998) (which adopted the narrower claim construction) because the written description in Laitram made clear that the asserted claims will bear only one interpretation. In Johnson Worldwide, there was no such unambiguous language in the claim; "nothing suggests that 'heading' is required to be the heading of a trolling motor." Johnson Worldwide, 175 F.3d at 991.

The facts of Johnson Worldwide are distinguishable from the Rambus patents here. The Johnson Worldwide court noted that the "many uses of the term throughout the . . . patent are consistent with a broader definition" and that the "[v]aried use of the term in the written description demonstrates the breadth of the term rather than providing a limited definition." Id. at 991. Thus, the dual usage of the term did not create "a special and particular definition" Id. Here, there are not varied uses of the term "bus," only a single multiplexed bus.

accustomed meaning of the term "frame" could be overridden by the inventor's explanation in the specification:

The parties agreed before the district court that the term "frame" can in general usage be applied to bit-mapped display systems as well as to character-based systems The disagreement was as to whether the term "frame" in the '669 claims embraced this general usage, or whether the term would be understood by persons of skill in this field as limited to the character-based systems described in the '669 patent.

Wang, 197 F.3d at 1381. As is true here, the only system described and enabled in the specification and drawings in Wang used the narrower, specific arrangement of the character-based system. Id. at 1382. The only time that the patent mentioned non-character-based systems was in the "Background of Invention" section. Id. The Federal Circuit agreed with the district court's conclusion that those references were merely acknowledgments of the state of the prior art, not an enlargement of the patent's invention; and that a person skilled in the field of art would not have understood that those references were included in the applicant's invention. Id. Similarly, Rambus is limited to the description set forth in the specification, which is only a description of the multiplexed bus. 

In an effort to distance the claims from the specification, Rambus argues that one skilled in the art would recognize that any kind of bus could be used with the many inventions of the specification, not just the new multiplexed bus. The Federal

Circuit has rejected this exact argument, which attempts to escape
the language of the specification. See Watts, 232 F.3d at 883
(inventor's arguments that "one of ordinary skill would be aware of
a myriad ways to effect a sealing connection . . . may be true,
[but] it does not overcome the fact that the specification
specifies that the invention uses misaligned taper angles"). The
fact that the inventions might conceivably be used with any kind of
bus does not overcome the oft-repeated assertions in the
specification which describe, and even tout, the new Rambus
inventive bus while demonstrating that the inventions are to be
used with the multiplexed bus.

Rambus next argues that Infineon is trying to improperly limit
the scope of the claim to the limitations described in the
preferred embodiment. See Karlin Tech., Inc. v. Surgical Dynamics,
Inc., 177 F.3d 968, 973 (Fed. Cir. 1999) ("The general rule, of
course, is that the claims of a patent are not limited to the
preferred embodiment, unless by their own language."); CVI/Beta
Ventures, 112 F.3d at 1158 ("as a general matter, the claims of a
patent are not limited by preferred embodiments"). The
specification, however, clearly distinguishes between the
"invention" of the multiplexed bus and "the preferred embodiment"
of the bus. The patent often describes the broader invention of a
bus multiplexed for address, data and control information. This
description is then followed by a narrower description of the

'preferred embodiment' that is an implementation of the multiplexed bus.

For example, the patent states:

The new bus includes clock signals, power and multiplexed address, data and control signals. In a preferred implementation, 8 bus data lines and an AddressValid bus line carry address, data and control information for memory addresses up to 40 bits wide. Persons skilled in the art will recognize that 16 bus data lines or other numbers of bus data lines can be used to implement the teaching of the invention.

'918 patent, col. 3, l. 67 through col. 4, l. 7. (emphasis added).

The new multiplexed bus is the broadly defined invention and the preferred embodiment has certain characteristics such as 8 or 16 multiplexed lines, an AddressValid line, and addresses of up to 40 bits. Numerous references in the specification highlight these differences.¹³ Usually, the preferred embodiments described in the

¹³ See e.g. 918 patent, col. 5, ll. 37-50 ("The bus carries substantially all address, data, and control information needed by devices for communication with other devices on the bus. . . . Using the organization described herein, very large addresses (40 bits in the preferred implementation) and large data blocks (1024 bytes) can be sent over a small number of bus lines 8 plus one control line in the preferred implementation."); '918 patent col. 14, ll. 49-67 ("In the bus-based system of this invention" a master can use the device ID to access a specific device "including the address and control registers. In the preferred embodiment, one master is assigned to carry out the entire system configuration process." (emphasis added); '918 patent, col. 16., ll. 12-21 ("The bus architecture of this invention can include more than one master device. The reset or initialization sequence should also include a determination of whether there are multiple masters on the bus, and if so to assign unique master ID numbers to each. Persons skilled in the art will recognize that there are many ways of doing this. For instance, the master could poll each device to determine

specification give a technical example of how the overall invention works, thus helping to explain the claim language. "Although claims are not necessarily restricted in scope to what is shown in a preferred embodiment, neither are the specifics of the preferred embodiment irrelevant to the correct meaning of claim limitations." Phonometrics, Inc. v. Northern Telecom, Inc., 133 F.3d 1459, 1466 (Fed. Cir. 1998). Where, as here, the several embodiments described in the specification each involves only a multiplexed bus, that weighs heavily in construing the term bus to mean a multiplexed bus. See Wang, 197 F.3d at 1383.

Finally, it is significant that the specification only mentions the generic (or "dedicated") bus architecture in the "Comparison of Prior Art" section.¹⁴ In these references to "bus," however, the inventors are distinguishing their new inventive bus from the prior art. The inventors explain that "[n]one of the buses described in patents or other literature use only bused connections. All contain some point-to-point connections on the backplane." '918 Patent, col. 2, l. 67 to col. 3, l. 3. Thus, it

what kind of device it is") (emphasis added).

¹⁴ For example, the specification explains that the bus of an earlier patent (U.S. Patent No. 3,821,715) "multiplexes addresses and data over a 4-bit wide bus and uses point-to-point control signals to select particular RAMs or ROMs." '918 patent, col. 2 ll. 13-15. The specification also explains that in the DRAM of a previous patent (U.S. Patent 4,449,207) "[t]he external interface to this DRAM is conventional, with separate control, address and data connections." '919 patent, col. 2, ll. 32-33.

does not help Rambus to point out, as it does, that this text of the comparison uses the same term ("bus") to describe a completely different architecture from the "new" bus which, according to Rambus, means that the term "bus" must necessarily encompass any set of information transfer lines, including those cited as prior art. The "Comparison With Prior Art" section states only what the invention does not cover; and, in so doing, the specification expressly distinguishes the prior art buses from the disclosed bus of the invention. Of course, it is settled that "[c]laims are not correctly construed to cover what was expressly disclaimed." Culter Corp. v. A.E. Staley Mfg. Co., 224 F.3d 1328, 1331 (Fed. Cir. 2000) (description in specification that distinguished other types of catalysts "effected a disclaimer of the other prior art acids"). See also Wang, 197 F.3d at 1382 (references to "bit-mapped" protocols in "Background of Invention" were acknowledgments of the state of the art and not an enlargement of the invention described in the patent). The argument which Rambus makes based on the term "bus" as used in the discussion of prior art runs afoul of this basic precept of claim construction.

In a further effort to use the discussions of prior art to support its proposed definition of "bus," Rambus relies on Clearstream Wastewater Sys. v. Hydro-Action, Inc., 206 F.3d 1440 (Fed. Cir. 2000), to argue that its inventions involve "combination

claims,"; therefore it is entirely permissible to include both the new and the generic buses in its inventions:

In construing the disputed claim limitations, it must be kept in mind that the claims at issue in this case are combination claims. Combination claims can consist of new combinations of old elements or combinations of new and old elements. Because old elements are part of these combination claims, claim limitations may, and often do, read on prior art.

Id. at 1445 (internal citations omitted). Further, Clearstream explains that:

Clearly, the written description does point out the disadvantages of the [prior art] rigid conduit system and the advantages of the [new] flexible-hose system. However, the written description does not require that only the new flexible-hose system, but not the old, rigid conduit system, could be used in the claimed wastewater treatment plant. It is well established in patent law that a claim may consist of all old elements . . . for it may be that the combination of old elements is novel and patentable. Similarly, it is well established that a claim may consist of all old elements and one new element, thereby being patentable.

Id. (emphasis added).

Infineon properly agrees that combination claims can include some, or even all, prior art elements. However, Infineon also is correct in asserting that the proper framework for the current analysis is whether one of ordinary skill in the art would understand the Rambus disclosure to assert the combination theory recently embraced by Rambus. The patent specification here does

not support that theory because, unlike Clearstream, the specification in Rambus' patents do not describe the generic prior art bus in combination with any of the claims. Indeed, the "Comparison With Prior Art" discussion is at considerable pain to dissociate the inventive bus, and its uses, from the prior art, and to establish a similar disconnect of the other inventions from the prior art. '918 Patent, col. 2, l. 7 to col. 3, l. 47.

Thus, the specification clearly demonstrates that when the inventors used the term "bus" in the claims, they were referring to the new multiplexed bus described in the specification.¹⁵ Upon reading the patent, one skilled in the art would conclude that the patentee explicitly defined bus: "[t]he present invention includes a memory subsystem comprising at least two semiconductor devices . . . connected in parallel to a bus where the bus includes a plurality of bus lines for carrying substantially all address, data and control information needed by said memory devices." '918 patent, col. 3, ll. 50-55. Nothing in the specification - no drawing and no embodiment - indicates that the bus in the claims has the dictionary definition that Rambus now asserts.

¹⁵ Other than the "Comparison with Prior Art" section, the patent specification only once indicates that a bus can be anything but the multiplexed bus: "Persons skilled in the art recognize that certain devices, such as CPUs, may be connected to other signals lines and possibly to independent buses, for example a bus to an independent cache memory, in addition to the bus of this invention." '918 patent, col. 5, ll. 54-57. In this reference, the inventors clearly distinguish between the multiplexed bus of the invention and any other kind of bus to be used in the system.

3. The File History

Despite the obvious descriptions in, and implications of, the patent itself, Rambus argues that the patent history teaches that the term "bus" includes more than just the multiplexed bus. While it is doubtful that a court should look to the patent history to contradict the unambiguous meaning described in the specification, see Multiform Desiccants, 133 F.3d at 1478 ("[w]hen the specification explains and defines a term used in the claims, without ambiguity or incompleteness, there is no need to search further for the meaning of the term"), the patent history here does not in any fashion clarify the scope of the disputed term. Rambus relies upon two statements made, and actions taken, in the prosecution of the patents stemming from the 1990 '898 application.

In June 1997, during the prosecution of the parent application¹⁶ to the '263 patent (which is also the grandparent to the '918 patent),¹⁷ the Patent Examiner issued a requirement for

¹⁶ This application eventually issued as U.S. Patent No. 5,841,580.

¹⁷ When considering a patent's prosecution history, it is proper to look to statements made in the prosecution of related patents stemming from the same application, as is the case here. See Elkay Mfg. Co. v. Ebco Mfg. Co., 192 F.3d 973, 980 (Fed. Cir. 1999) ("When multiple patents derive from the same initial application, the prosecution history regarding a claim limitation in any patent that has issued applies with equal force to subsequently issued patents that contain the same claim limitation."); Mark I Marketing Corp. v. R.R. Donnelley & Sons Co., 66 F.3d 285, 291 (Fed. Cir. 1995) ("Thus, the relevant prosecution history here includes not only the '659 application but also the parent '815 and grandparent '668 applications."); Jonsson v.

restriction under 35 U.S.C. § 121,¹⁸ finding that this patent claimed two distinct inventions. The examiner divided the claims into two groups, one group describing a plurality of conductors to be used with the multiplexed bus and the second group describing an access-time register within the memory device (the latency invention).¹⁹ Asserting that the groups were not "connected in

Stanley Works, 903 F.2d 812, 818 (Fed. Cir. 1990) (prosecution history of parent application is relevant to understanding scope of claims issuing in a continuation-in-part application).

¹⁸ A requirement for restriction is issued by the PTO when a patent application contains more than one distinctly claimed invention. 35 U.S.C. § 121.

¹⁹ The June 9, 1997 Office Action explains:

4. Restriction to one of the following inventions is required under 35 U.S.C. 121:

Group I. Claims 151-55, drawn to a memory device having a plurality of conductors being multiplexed for sequentially receiving an address, classified in Class 365, subclass 230.02.

Group II. Claims 156-158, drawn to a semiconductor device having at least one access-time register, classified in Class 395, subclass 290.

The inventions are distinct, each from the other because of the following reasons:

5. Inventions I and II are disclosed as different combinations which are not connected in design, operation or effect. These combinations are independent if it can be shown that (1) they are not disclosed as capable of use together, (2) they have different modes of operation, (3) they have different functions, or (4) they have different effects. (MPEP 806.04, MPEP 808.01). In the instant case the combinations

design, 'operation, or effect," the Patent Examiner required the inventors to elect to pursue only one group of claims. Rambus prosecuted the claims in Group II (the latency invention), resulting in the '580 patent.

From this action by the PTO and from the fact that the same Patent Examiner reviewed the '263 and '918 patents in suit, Rambus asks the Court to make the leap in logic that the PTO must have understood that the multiplexed bus was not necessary for every other invention arising from the specification. This kind of speculation into the motivations of the patent examiner is not useful to a reviewing court or a competitor reading the patent history. "It is the applicant's representations during the prosecution that potentially shed light on the construction of the claims." Laitram Corp. v. Morehouse Indus., Inc., 143 F.3d 1456, 1462-63 (Fed. Cir. 1998) (emphasis in original) (rejecting argument

[sic] the memory device in Group I does not require the access-time register of Group II, and the semiconductor device in Group II does not require the plurality of conductor being multiplexed to receive an address as claimed in Group I.

6. Because these inventions are distinct for the reasons given above and have acquired a separate status in the art as shown by their different classification, and the search required for invention I is not required for invention II, restriction for examination purposes as indicated is proper.

(emphasis added).

that meaning could be derived from the representation of the reexamination requester). The snippet of patent history upon which Rambus relies only shows that: (a) a single Patent Examiner at one time indicated that some claims should not be lumped together, and (b) that, rather than making an affirmative response to this restriction, Rambus chose to drop the claims for the multiplexed bus and pursue the latency invention. No more can be inferred from this exchange.

Rambus also relies on a second piece of evidence contained in the file history. Specifically, in November 1995, a different Patent Examiner rejected claims pending in the grandparent to the '804 patent as being obvious in view of prior art reference, U.S. Patent No. 5,129,069 to Helm, et al. Under Rambus' view of the file history, the Patent Examiner must have equated the generic term "bus" (recited in claims 176-181 of the grandparent application) with the non-multiplexed bus contained in the Helm patent when he initially rejected the claims. Nevertheless, this same Patent Examiner allowed claim 26 of the '804 patent (which contains a reference to an external bus) to issue without requiring that the term "bus" be limited to a multiplexed bus. Again, this kind of guessing as to what a Patent Examiner may have been thinking is not generally helpful to construing the claim terms because it requires both the court and the public to pour over oftentimes complex and voluminous patent histories, speculate as to

the motivation behind an office action, and then follow the patents in an effort to divine whether that same Patent Examiner may have had reason to construe another claim in the same manner. This invitation to haphazard guesswork certainly cannot be considered sufficiently reliable to trump the clear language of the specification. See Vitronics, 90 F.3d at 1582 ("Usually, [the specification] is dispositive; it is the single best guide to the meaning of a disputed term").

Moreover, the standard for construing claims in the patent application process is far different than the standard for construing claims in a litigation context. Patent examiners construe claims under a broader standard than that used by a court in undertaking claim construction. The Federal Circuit has held that "[i]t would be inconsistent with the role assigned to the PTO in issuing a patent to require it to interpret claims in the same manner as judges who, post-issuance, operate under the assumption that the patent is valid." In re Morris, 127 F.3d 1048, 1054 (Fed. Cir. 1997). In the posture of a claim construction during litigation, if the intrinsic evidence is ambiguous, "another claim construction canon comes into play. Because the applicant has the burden to 'particularly point[] out and distinctly claim[] the subject matter which the applicant regards as his invention' 35 U.S.C. § 112, ¶ 2 (1994), if the claim is susceptible to a broader and a narrower meaning, and the narrower one is clearly supported

by the intrinsic evidence while the broader one raises questions of enablement under § 112, ¶ 1, we will adopt the narrower of the two." Digital Biometrics, 149 F.3d at 1344. See also Athletic Alternatives, Inc. v. Prince Mfg. Inc., 73 F.3d 1573, 1581 (Fed. Cir. 1996) ("Were we to allow AAI successfully to assert the broader of the two senses of 'between' against Prince, we would undermine the fair notice function of the requirement that the patentee distinctly claim the subject matter disclosed in the patent from which he can exclude others temporarily.") Therefore, even if one were to conclude that the patent history casts doubt on the clear meaning of the specification (which it does not), Rambus should be limited to the embodiment and description of a multiplexed bus set forth in the specification because it is the narrower of the two constructions.

The simple fact here is that reference to the file history does not contradict the clarity given by the specification. What Rambus has done is fixate upon two isolated events in the file history and, without connecting them to the issued patents, urges the Court to ascribe significance to the events by divining what an examiner must have meant by directing a certain action. That kind of sophistry is not among the tools available for claim construction under the carefully defined protocol established for that task by the Federal Circuit.

4. Claim Differentiation

Rambus relies on the doctrine of claim differentiation to support its contention that "bus" means only a "generic" bus. The doctrine presumes "a difference in meaning and scope when different words or phrases are used in separate claims. To the extent that the absence of such difference in meaning and scope would make a claim superfluous, the doctrine of claim differentiation states the presumption that the difference between claims is significant."

Toro Co., 199 F.3d at 1302 (Fed. Cir. 1999) (quoting Tandon Corp. v. United States Int'l Trade Comm'n, 831 F.2d 1017, 1023 (Fed. Cir. 1987)). Rambus highlights the claims contained in Rambus U.S. Patent No. 5,983,320 (the '320 patent),²⁰ in which the independent claims of the '320 patent cover the concept of the "new" multiplexed bus. Each claim contains qualifying language to limit the bus to one that carries multiplexed address, data and control information over the same bus. For example, claim 7 of that patent claims a method for programming memory having a bus, where the bus "compris[es] a group of general purpose signal lines carrying substantially all of the time-division multiplexed address, data and control information for a memory transaction, wherein the

²⁰ The '320 patent is a "sister" or "brother" patent to the '804 patent. Rambus bases the notion of cross-patent claim differentiation on footnote 2 of Laitram Corp. v. Morehouse Indus., Inc., 143 F.3d 1456, 1460 n.2 (1998). Because the argument of claim differentiation fails for other reasons, it is assumed, without deciding, that this is a proper use of related patents and their prosecutions.

address information is indicative of a range of addresses for a corresponding one of the individually addressable discrete memory sections of the memory device . . ."²¹ Therefore, Rambus argues that, when it wanted to limit the term "bus" to a bus that carries multiplexed information, it knew how to do so.

This argument is unavailing because it too contradicts the clear meaning of the specification. "The doctrine of claim differentiation cannot broaden claims beyond the scope that is supported by the specification." ATD Corp. v. Lydall, Inc., 159 F.3d 534, 541 (Fed. Cir. 1998). See also Multiform Desiccants, 133 F.3d at 1480 (same). "Although the doctrine of claim differentiation may at times be controlling, construction of claims is not based solely upon the language of other claims; the doctrine cannot alter a definition that is otherwise clear from the claim language, description, and prosecution history." O.I. Corp. v. Tekmar Co., Inc., 115 F.3d 1576, 1582 (Fed. Cir. 1997) (concluding "that the description provides a clear meaning for the language of the claim in this case and that it trumps the doctrine of claim differentiation"). See Toro Co., 199 F.3d at 1302 (claim differentiation "does not override clear statements of scope in the

²¹ According to Rambus, other Rambus patents include similar limiting language: U.S. Patent No. 5,995,443, Claim 33 ("the bus further includes a plurality of conductors terminated by an impedance to a power source") and U.S. Patent No. 6,032,215, Claims 33 and 37 (same) and Claim 38 ("the bus further includes a plurality of conductors wherein each conductor is terminated at an end by a resistor to a power terminal.")

specification and the prosecution history"). "The presumption that separate claims have different scope 'is a guide, not a rigid rule.'" ATD Corp., 159 F.3d at 541 (quoting Autoapiro Co. of Am. v. United States, 384 F.2d 391, 404 (1967)). Having determined that the written specification limits the term "bus" to a multiplexed bus, it would be impermissible to allow Rambus to rely upon claim differentiation (citing to other patents) to broaden the meaning of the term.²²

5. The Extrinsic Evidence

A review of the intrinsic evidence clearly demonstrates that when the term "bus" is used in the claims, it means the new inventive Rambus multiplexed bus. "Because the intrinsic record is

²² Along these same lines, Rambus argues that the claims of the original '898 application specifically claim a multiplexed bus, therefore the reasoning behind claim differentiation would apply to give "bus" a generic meaning within the specification because the original claims are part of the specification. See In re Dossel, 115 F.3d 942, 945 (Fed. Cir. 1997) ("The statute thus makes clear that under current law the specification of a patent consists of, and contains, both a written description of the invention and the claims."); Northern Telecom, Inc. v. Datapoint Corp., 908 F.2d 931, 938 (Fed. Cir. 1990) ("The original claims as filed are part of the patent specification.") However, Dossel explains that "[m]odern usage . . . does not always conform to that statutory structure. For example, when discussing the process of claim construction, it is not uncommon for the process to be described as requiring an examination of the claims, the specification, and the prosecution history, treating them as distinct entities." 115 F.3d at 945. To the extent that the claims of the original application (which never issued), indicates that the inventors distinguished between the new multiplexed bus and a generic bus, this difference does not trump the clear descriptions in and implications of the written description.

clear, [the court] do[es] not give weight to an inconsistent dictionary definition," Digital Biometrics, 149 F.3d at 1346, as offered by Rambus.

A reading of the entire specification, without parsing it into individual quotes, unmistakably conveys that one of the primary benefits of every invention claimed in the patents in suit, and described in the specification, is to increase the speed of operation of the memory device. High speed access is the crown jewel of the specification and to that end, the multiplexed bus, in combination with the other inventions, increases the transfer speeds and decreases the amount of space occupied by the transfer lines. One skilled in the art reading the specification would certainly conclude that the "bus" meant to be used with the inventions is the new, inventive, high-speed, multiplexed bus. This conclusion is supported by the testimony of Joseph McAlexander, an expert who is experienced in the art and whose explanation for reaching that conclusion is highly credible because it is fully consonant with the specification and the claim language as explicated by the specification.²³ Mr. McAlexander's testimony

²³ See Markman Hearing, Tr. p. 379, l. 20 to p. 380, l. 3 (Mr. McAlexander states, "Because the patent very strongly distinguishes numerous times the multiplex bus of the invention from the prior art, and states specifically in numbers of places, that the bus architecture of this new bus design is essential for the type of high speed activity that is required across the bus, and it distinguishes from the prior art because the prior art is stated not to be able to give that high speed type of transaction.")

is consistent with, and complimentary of, the intrinsic evidence.²⁴ On the other hand, the testimony of Rambus' expert, Dr. William Huber, is at odds with the intrinsic evidence and depends on a dictionary definition (other extrinsic evidence) that is not consistent with the many descriptions given by the inventors in the specification.

6. Construction

For the foregoing reasons, the term "bus" means a multiplexed set of signal lines used to transmit address, data and control information.

B. "Block Size Information"

In Rambus' invention, the user can specify the amount of data to be transferred over the bus during a bus transaction. This value is represented by the term "block size." The parties differ as to exactly how this value is to be measured. Rambus argues that "block size" is "the number of sequential data bits to be read from or written to the memory." In essence, Rambus reads block size to be a function of the number of sequential transactions on a bus necessary to respond to the transaction request. Infineon posits that "block size" "specifies the total amount of data that is to be

²⁴ Other portions of Mr. McAlexander's testimony generally support this construction. See Markman Hearing, Tr., pg. 361, ll. 14-24; pg. 364, l. 22 to pg. 365, l. 22; pg. 367, l. 17 to pg. 368, l. 13; pg. 371, ll. 3-19.

transferred on the bus in response to a transaction request." In other words, Infineon measures block size as a function of size or the amount of data to be transferred over the bus.

1. The Claim Language

The term "block size" occurs numerous times throughout the claims of the '918 patent and the '214 patent.²⁵ Most of the claims indicate that block size information defines the amount of data to be output or input by the memory device.²⁶ Indeed, there is nothing in the text of any claim which employs the term "block size" to indicate that "block size" means anything other than the amount of data to be transferred on the bus in response to some sort of transaction request. Thus, from reading the language of the claims, one skilled in the art would conclude that block size information is an instruction indicating the amount of data to be output (or input) by the memory device.

²⁵ See claims 1, 2, 6, 9, 13, 15, 16, 18, 19, 20, 24, 29, 30, 31, 33, and 34 of the '918 patent and claims 1, 4, 6, 9, 10, 15, 16, 18, 21 and 25 of the '214 patent.

²⁶ See e.g. '918 patent, claim 1 ("first block size information defines a first amount of data to be output by the memory device. . . ."); '918 patent, claim 6 ("the memory device outputs the first amount of data corresponding to the first block size information. . . ."); '918 patent claim 13 ("the first block size information is a binary representation of the amount of data to be output after receipt of the first read request."); '214 patent, claim 1 ("first block size information defines a first amount of data to be output onto a bus. . . ."); '214 patent, claim 6 ("first block size information is a binary code indicative of the first amount of data to be output in response to the read request").

Rambus nonetheless urges that block size means "the number of sequential data bits to be read from or written to the memory" (presumably in response to a transaction request). Rambus does not identify any aspect of the claim language that would support its preferred definition. Additionally, it is worth noting that nothing in the claim language has been cited, or for that matter argued, as supporting the temporal or order requirements which would inhere in a sequential-based definition.

2. The Specification

In general terms, the specification explains that "[o]ne object of the present invention is to use a new bus interface built into semiconductor devices to support high-speed access to large blocks of data from a single memory device by an external user of the data, such as a microprocessor, in an efficient and cost-effective manner." '918 patent, col. 3, lines 21-25 (emphasis added). See also '918 patent, col. 4, lines 15-16 ("The bus supports large data block transfers . . .").

In discussing the preferred method of Device Address Mapping and the address registers therein employed, the specification explains that:

The address registers can include a single pointer, usually pointing to a block of known size, a pointer and a fixed or variable block size value or two pointers, one pointing to the beginning and one to the end (or to the "top" and "bottom" of each memory block.

'918 Patent, col. 7, ll. 36-41. This text and that which follows it clearly bespeaks volume or amount as the measure of a block, not sequence or time.

The several other references to block size in the specification also teach that the term relates to amount of data not the order and timing of bits of data in a particular sequence.²⁷

3. The Extrinsic Evidence

The construction offered by Rambus purports to be grounded in a table in column 11 of the specification. The table is in a section which refers to the preferred embodiment detailed in Figure 4 of the patent.²⁸ The specification states:

BlockSize[0:3] specifies the size of the data block transfer. If BlockSize[0] is 0, the remaining bits are the binary representation of the block size (0-7). If BlockSize[0] is 1, then the remaining bits give the block size as a binary power of 2, from 8 to 1024. A zero-length block can be interpreted as a special command, for example, to refresh a DRAM without returning any data, or to change the DRAM from page mode to normal access mode or vice-versa.

²⁷ '918 Patent, col. 11, ll. 1-5, ll. 41-48; col. 16, ll. 26-35, ll. 44-47; col. 17, ll. 1-2; col. 20, ll. 18-22.

²⁸ Figure 4 is replicated later in the Memorandum Opinion at II.C.2.

BlockSize[0:2]	Number of Bytes in Block
0-7	0-7 respectively
6	8
9	16
10	32
11	64
12	128
13	256
14	512
15	1024

Persons skilled in the art will recognize that other block size encoding schemes or values can be used.

'918 patent, col. 11, ll. 41-63.²⁹

Rambus uses Table 11, as interpreted by its expert, Dr. Huber, as the basis for its construction that block size information is the number of sequential transfers necessary to carry the desired information over the bus line. The table is but one of many encoding schemes and does not purport to define or explain the meaning of block size generally. According to Dr. Huber, the block size indicated in the chart corresponds to the number of sequential transfers necessary to output the data onto the preferred embodiment 8-line bus. That is certainly not apparent from the patent document. Furthermore, during the Markman hearing, Dr. Huber connected almost all of his opinions, not to the patent specification, but rather to a prepared animation demonstrating how block size should be measured as sequential transfers of data. The reason for such reliance seems quite clear -- there simply is no

²⁹ During the Markman hearing, the parties agreed that the "6" contained in the second row of the chart is incorrect due to an apparent copying error. This number should be an "8."

support in the patent document. Also, that view is flatly contradicted by Infineon's expert who explained that block size contains information specifying the total amount of data that is to be transferred. See Markman Hearing, Tr. pg. 442, ll. 11-19. Mr. McAlexander has testified that "[t]he person of ordinary skill in the art would come to [the conclusion] that block size . . . means amount, and in just a plain, simple ordinary meaning of size is a[n] amount, it's not when or how." Markman Hearing, Tr. pg. 439, ll. 11-16 (testimony of Mr. McAlexander).

4. Construction

Infineon's construction is grounded in the specification and the claim language because both sources of information rather clearly reflect that block size is an amount of data, not the order in which it is delivered. There is nothing in the specification to support Rambus' somewhat contorted definition of block size. It simply defies reason (and the specification) to conceive of size as a measure of time.

Moreover, the construction urged by Rambus utterly ignores the clear language of the claim that block size is associated with a transaction request. (See, e.g., '918 Patent, Claim 1-7 and all other claims (8 through 38) dependent upon Claim 1-7). Infineon's definition encompasses this connection and, for that additional reason, it is the definition that is necessitated by the claim language and by the specification. Thus, "block size" is construed

to mean, "information that specifies the total amount of data that is to be transferred on the bus in response to a transaction request."

C. "Read Request," "Write Request," and "Transaction Request"

The next three disputed terms are closely related, and, as the parties agree, it is appropriate to accord joint consideration. The purpose of memory devices (*i.e.*, a DRAM) is to store data for later use. To this end, when a controller (or master) accesses the memory device to either store or retrieve data, it must send that device an instruction indicating what type of transaction is to be performed. At the most basic level, a read request is an instruction to the memory device to read data from the memory cells; a write request instructs the memory device to write data to the memory cells; and a transaction request instructs the memory device to perform some function, which could include reading or writing data. The controversy surrounding these terms involves whether these requests must contain not only the instruction of what action to perform (found in an "AccessType" field), but also must include address information indicating where in the memory cells the data should be read or written. Rambus contends that read, write and transaction requests contain only the instruction of what action to perform. For example, it proposes that read request be defined as "an instruction to read data from specified

memory cell(s) of the memory. This instruction is specified by a binary code³⁰ provided to the memory device during a single clock cycle and received by the memory device in response to a clock transition."

Infineon, on the other hand, argues that such requests must contain both the instruction of what kind of action to perform and address information indicating where that action is to occur on the memory device. Address information, containing both row and column identifiers, tells the memory device where the desired data is located (or to be located) within the plurality of the memory cells. In Infineon's construction, a "read request" means "a series of bits transmitted over the bus that contain multiplexed address and control information needed to request a read of data from a memory device." In addition to objecting to the failure of Rambus' definition to include address information, Infineon disagrees with the limitations inherent in Rambus' definition, specifically the requirements that the request be a "binary code," that it can be no longer than "a single clock cycle" of information, and that it must be "received in response to a clock transition." In its view, these limitations are not required by

³⁰ Binary code, a term not in dispute here, is "a code that makes use of members of an alphabet containing exactly two characters, usually 0 and 1." *IEEE Standard Dictionary of Electrical and Electronics Terms*, 4th Ed., IEEE, Inc. NY, 4th Ed. Pg. 95.

the intrinsic evidence, and in some cases are actually inconsistent with the embodiments disclosed in the specification.

1. The Claim Language

Both parties agree that the terms "read request," "write request," and "transaction requests" are not terms of art and were used for the first time in the 1990 '898 application. Therefore, there is no ordinary and accustomed meaning for these terms. Some information, however, can be gleaned from the language of the claims.³¹

Infineon uses the language of claim 1 of the '918 patent, claim 1 of the '214 patent, claim 1 of the '263 patent and claim 26 of the '804 patent to demonstrate that all claims require that a device respond to 'read request':

1. A method of controlling a synchronous memory device, wherein the memory device includes a plurality of memory cells, the method of controlling the memory device comprises:

providing first block size information to the memory device, wherein the first block size information defines a first amount of data to

³¹ The term "read request" occurs in claims 1, 6, 8, 13, 18, 19, 24, 29, and 34 of the '918 patent; claims 1, 2, 6, 14, 15, 16, 18, and 29 of the '214 patent; claims 1, 2, 14, 24, 15, 27, and 30 of the '263 patent; and claim 26 of the '804 patent.

The term "write request" occurs in claims 2 and 20 of the '918 patent.

The term "transaction request" occurs in claims 18 and 25 of the '263 patent.

be output by the memory device onto a bus in response to a read request . . .

'918 patent, claim 1 (emphasis added).

1. A method of operating a synchronous memory device, wherein the memory device includes a plurality of memory cells, the method comprising:

providing first block size information to the memory device, wherein the first block size information defines a first amount of data to be output onto a bus in response to a read request . . .

'214 patent, claim 1 (emphasis added).

1. A synchronous semiconductor memory device having at least one memory section which includes a plurality of memory cells, the memory device comprises:

a programmable register to store a value which is representative of a delay time after which the memory device responds to a read request.

'263 patent, claim 1 (emphasis added).

26. An integrated circuit device having at least one memory section which includes a plurality of memory cells, wherein the integrated circuit device outputs data on an external bus synchronously with respect to first and second external clock signals, the integrated circuit device comprises:

a first internal register to store a value which is representative of a number of clock cycles to transpire before the integrated circuit device responds to a read request. . .

'804 patent, claim 26 (emphasis added).

The claims clearly so provide. Indeed, all but two of the disputed claims containing these terms³² explicitly state that information is supplied "in response to" a read request, a write request or other transaction request. It is, of course, true, as Infineon contends, that, in order to "respond" to a request (i.e. outputting or inputting data), the desired response can only occur if the selected device is given the information necessary to generate that response. Because one of ordinary skill in the art would understand that both address and control information are required for the memory device to respond to a request, that request must contain more than just the binary code or "AccessType" suggested by Rambus. Given the nature of the information and the way the invention works, it seems self-evident from the claim language that a request is, as Infineon posits, a series of bits transmitted over the bus containing address and control information. This conclusion is further buttressed by the explanations of Mr. McAlexander. See Markman Hearing, Tr. pg. 417, ll. 18-25 ("certainly the read has to have some control. It tells you what kind of transaction is being requested. If the memory is to respond to that, it must know where to respond from, what

³² Indeed, of those claims mentioning read request, write request and transaction request, only claims 14 and 29 of the '214 patent do not explicitly mention that the memory device is to respond to the read request and even those mention a read request in such a way as to indicate that the term means there what it means elsewhere (Claim 14 "before executing another read request"); (Claim 29 "after executing another read request.")

address"). Thus, the claim language, although not dispositive, strongly supports the view of read, write and transaction request taken by Infineon.

2. The Specification

Although not discussed as extensively as other terms, such as "bus," the terms here at issue are the subject of explication in the specification. For example, in the "Comparison With Prior Art" section, the inventors explain:

Yet another object of this invention is to provide a method for transferring address, data and control information over a relatively narrow bus and to provide a method of bus arbitration when multiple devices seek to use the bus simultaneously.

'918 Patent, col. 3, ll. 35-39 (emphasis added). This statement of object remarks the key role of address, data and control information. And, as explained in the cited text, and above in construing the term "bus," the significance of the invention of the system is to accomplish quickly the commands necessary to initiate a request and secure a response.

Then, in the ensuing "Summary of Invention" discussion, the inventors say that "[i]n this system of the invention, DRAMs and other devices receive address and control information over the bus and transmit or receive requested data over the same bus." '918 Patent, col. 4, ll. 9-11 (emphasis added). Though not explicitly mentioning a "request," this quote lends credence to the basic

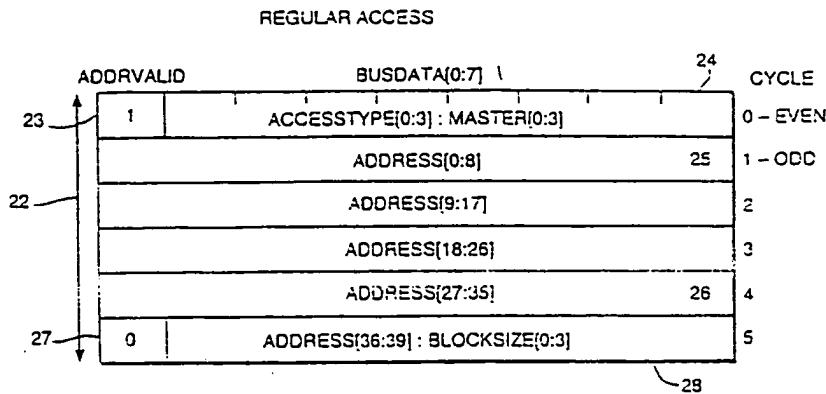
notion that a memory device should receive both address and control information in order to be able to transmit or receive data.

Then, shortly thereafter, in describing a preferred implementation of the invention, the specification explains how a bus transaction is initiated "by sending a request packet (a sequence of bytes comprising address and control information)." '918 Patent, col. 6, ll. 60-63). This, too, teaches that a request (be it a read request, write request, or transaction request) includes the address and control information necessary to accomplish the request.³³

Having established that the specification contemplates both address and control information are needed for a response, it is necessary to ascertain whether a read request should contain both categories of information. The definition which Rambus presses proposes that a transaction request would consist only of the "AccessType" found in the top row³⁴ of Figure 4, the preferred embodiment.

³³ Again, this conclusion is supported by the expert testimony of Infineon's Mr. McAlexander. See Trans. pg. 417, lines 20-22 ("I found the control and address information were required in every instance that it was addressed in the specification").

³⁴ As shown in Figure 4, the rows represent time or clock cycles.



FIG_4

In the preferred embodiment, the AccessType instruction, which contains the control information specifying the type of request, would be a binary code 4-bits wide. The specification explains:

The AccessType filed [sic: field] specifies whether the requested operation is a read or write In a preferred implementation, AccessType[0] is a Read/Write switch: if it is a 1, then the operation calls for a read from the slave (the slave to read the requested memory block and drive the memory contents onto the bus); if it is a 0, the operation calls for a write into the slave (the slave to read data from the bus and write it to memory).

'918 patent, col. 9, ll. 47-56. Rambus would limit the terms "read request," "write request" and transaction request" to only this AccessType field. In contrast, according to Infineon's interpretation, the requests must contain both the AccessType control information and the address information indicated on the remaining rows of Figure 4. To support its definition, Infineon points to the following passage from the specification:

In a preferred implementation of the invention, to initiate a bus transfer over the bus, a master sends out a request packet, a contiguous series of bytes containing address and control information . . .

The device-selection function is handled using the bus data lines. AddrValid is driven, which instructs all slaves to decode the request packet address, determine whether they contain the requested address, and if they do, provide the data back to the master (in the case of a read request) or accept data from the master (in the case of a write request) in a data block transfer.

'918 patent, col. 8 l. 66 through col. 9, l. 4. The specification also explains that "[i]n some cases, a slave [memory device] may not be able to respond correctly to a request, e.g. for a read or write. In that situation, the slave should return an error message . . . or a retry message." '918 patent, col. 12, ll. 4-8. These references and others³⁵ illustrate that the memory should respond to the request. In order to respond, the memory device also must be given address information specifying where the data is to be read or written.

Additionally, though one must understand the technology to comprehend the import of the statement, the specification actually states that address rows are to be accessed during a request. The patent explains that the DRAM sense amps should be pre-charged and

³⁵ See e.g. '918 patent, col. 8, ll. 48-49 ("a slave should preferable respond to a request in a specified time"); '918 patent, col. 8, ll. 24-29 ("AddrValid is used to indicate when the bus is holding a valid address request, and instructs a slave to decode the bus data as an address and, if the address is included on that slave, to handle the pending request.")

"[t]his 'precharging allows access to a row in the RAM to begin as soon as the access request for either inputs (writes) or outputs (reads) is received and allows the column sense amps to sense data quickly." '918 Patent, col. 10, ll. 21-24 (emphasis added). Because one skilled in the art would recognize that "row" refers to a particular location on the plurality of memory cells, it follows that address information must be conveyed in order to access that row. That address information is contained in "the access request for either inputs (writes) or outputs (reads)."

Rambus definition would only indicate what type of operation to take place. See Markman Hearing, Tr. pg. 122, ll. 11-13 (testimony of Dr. Huber) ("We don't need the rest of the information [address information] to know that it's a read request"). The specification and the claims, however, clearly demonstrate the memory devices are not only to recognize the requested operation, but also respond to the request. Even Dr. Huber admitted that address information must be received for there to be a response. Dr. Huber took the view that this information could be conveyed at some other unspecified time. See Markman Hearing, Tr. pg. 141, ll. 20-22 (testimony of Dr. Huber). That approach is untenable because nowhere in the specification is it mentioned that address information should be sent at any other time than contemporaneous with the request.

Nor does the specification support the other foundational components of the narrow view of these terms expressed by Rambus. The construction urged by Rambus essentially attempts to equate the term "read request" with "AccessType," as shown in Figure 4, as the predicate for its requirement that "a read request" must be a binary code, occur in a single clock cycle, and be in response to a clock transition. The specification offers no warrant for such a limited construction, and, as Infineon points out, Rambus here is attempting artificially to limit the invention to the preferred embodiment of Figure 4, which describes a bus transaction that uses the preferred implementation of a 9-bit wide external bus. See '918 patent, col. 9, ll. 26-27 ("Each request packet uses all nine bits of the multiplexed data/address lines"). Because Figure 4 indicates that the AccessType is only 4-bits wide, it is possible for Rambus' proposed definition to occur in a single clock cycle. However, this requirement stems solely from Rambus' view that AccessType is a request. If a request contains both control and address information, then this would not be true. See Markman Hearing, Tr. pg. 431, ll. 8-12 (testimony of McAlexander) ("There is a specific control set of bits called the access type that does occur as a set of bits in a particular single cycle as shown in a preferred embodiment, but all that does is establish the type.") Similarly, there is nothing in the specification to support Rambus' requirement that the transaction request be received by the memory

device in response to a clock transition. See Markman Hearing, Tr. pg. 432, ll. 2-9 (testimony of McAlexander) (indicating that nothing in the specification supports this requirement).

Rambus contends that Infineon's construction incorrectly equates "read request" (or "write request" or "transaction request") with a request packet, arguing instead that a read request is actually a component of a request packet. For example, the specification states that "FIG. 4 shows the format of a request packet." '918 patent, col. 4, l. 66. It also explains that, in a preferred implementation, "a master sends out a request packet, a contiguous series of bytes containing address and control information." '918 patent, col. 8, ll. 60-63. The criticism is superficially appealing; however, the confusion results in large part from the fact that the specification uses the term "request" and "request packet" interchangeably. For example, the inventors explain, "FIG. 5 illustrates the format of a retry message 28 which is useful for read requests, . . . All DRAMs and masters can easily recognize such packet as an invalid request packet, and therefore a retry message." '918 patent, col. 12, ll. 33-39. See also '918 patent, col. 12, ll. 49-52 ("The master sends request packets and keeps track of periods when the bus will be busy in response to that packet. The master can schedule multiple requests so that the corresponding data block transfers do not overlap."); '918 patent, col. 12, ll. 58-61 ("Situations will arise, however,

where two or more masters send a request packet at or about the same time and the multiple requests must be detected. . . .) That drafting lapse is unfortunate, but it certainly is not dispositive because that text too must be interpreted in perspective of the whole specification.

Considering the claim language and the specification in its entirety and for the reasons explained above, the construction offered by Infineon is better supported by the patent document. Although that construction results in some overlap in the meanings of request and request packet, that overlap is inherent in the patent specification itself. Indeed, the most significant passage of the specification discussing read requests and write requests indicates that such a request is related to (if not synonymous with) a request packet: "AddrValid is driven, which instructs all slaves to decode the packet address determine whether they contain the requested address, and if they do, provide the data back to the master (in the case of read request) or accept data from the master (in the case of a write request) in a data block transfer." '918 Patent, col. 8, l. 66 through col. 9, l. 4 (emphasis added).

3. The File History

Those constructions derived from the claim language and specification are supported by the fact that, in the prosecution of the '804 patent, Rambus made statements to the PTO relating to the term "transaction request." In February 1999, Rambus submitted a

Preliminary Amendment in U.S. Patent App. 08/798,525 (issued as the '804 patent) in which it admitted that transaction requests are not simply a single clock-cycle access-time code. In response to a rejection by the Patent Examiner, Rambus stated that a "transaction request" contains identification information:

When the identification information contained in the transaction request corresponds to the identification value stored in the internal register in a particular memory device on the module, that memory device executes the transaction request. Memory devices on the module having identification values which do not correspond to the identification information contained in the transaction request do not execute or respond to the request.

Supplemental Preliminary Amendment, U. S. Patent App. 08/798,525, p. 12 (emphasis added). Thus, Rambus explicitly represented that a transaction request contains more than just a binary code in the AccessType field: the above passage shows that device identification information also is contained in the transaction request. While this representation does not necessarily imply that Infineon's definition is unquestionably correct, it certainly undermines the construction now urged by Rambus.

4. Claim Differentiation

Claim 15 of the '214 patent refers to a "read request" without further limitation, while dependent claim 22 recites: "The method of claim 15 wherein the first block size information and the first read request are contained in a request packet." Rambus argues

that this language distinguishes a "read request" from a "request packet." As stated previously, the specification sometimes uses the terms "request" and a "request packet" interchangeably. Notwithstanding that drafting laxity, the differences in claim 15 and 22 do not refute the notion that a request contains address and control information. These claims simply add a third type of information, block size information, as a component of a request packet.

Given that the claim language clearly illustrates that a memory device is to respond to a read, write or transaction request and that Rambus has not explained how the device would respond without receiving address, data and control information, the claim language on its face supports the requirement that requests contain both address and control information. The specification, while not pellucid, also indicates that a request must contain such information so that it can respond to the request, whether the request be packetized or not. Rambus' narrow definition is not supported by the specification, and indeed, is refuted by the file history. Therefore, it is appropriate to conclude that "read request," "write request," and "transaction request" contain both address and control information indicating what type of transaction to perform and where the data should be located on the memory device.

5. Construction

For the foregoing reasons, the term "read request" is construed to mean "a series of bits transmitted over the bus that contain multiplexed address and control information needed to request a read of data from a memory device." The term "write request" is construed to mean "a series of bits transmitted over the bus that contain multiplexed address and control information needed to request a write of data to a memory device." The term "transaction request" is construed to mean "a series of bits transmitted over the bus that contain multiplexed address and control information needed to perform a transaction over the bus with a memory device."

D. "First and Second External Clock Signals"

The parties agree that the bus of the invention carries two external clock signals³⁶ which pace the exchange of information over the bus and provide timing synchronization for the memory system. The dispute arises over whether the second external clock signal must contain information that is different from the timing information sent by the first clock signal.

Although referred to as a "clock" by one skilled in the art, the clock of a memory chip is actually a set of timing information

³⁶ As has been the convention of both the patent documents and the parties, the terms "clock signal" and "clock" are used interchangeably.

derived from an oscillating reference voltage ("V_{REF}") which cycles between two voltage levels. Rambus' proposed definitions do not require that the two signals contain different timing information, while Infineon's proposed definitions require that the second signal contain different information from the first.

1. The Claim Language

Every asserted claim in the '214 patent (the double data rate invention) and the '804 patent (the delayed lock loop invention) contains the terms "first external clock signal" and "second external clock signal."³⁷ Most of the claims simply indicate that data is to be output on the bus in response to the first and second external clock signals. Three claims, however, reveal that the two clock signals can be used by the memory device to create an internal clock:

25. The method of claim 15 further including generating at least one internal clock signal using the first and second external clock signals wherein the first amount of data corresponding to the first block size information is output onto the bus synchronously with respect to at least one internal clock signal.

'214 patent, claim 25.

26. The method of claim 25 further including generating a first internal clock signal using a delay locked loop and the first and second external clock signals.

³⁷ See claims 1, 2, 4, 9, 10, 11, 15, 16, 18, 24, 25, and 26 of the '214 patent and claim 26 of the '804 patent.

'214 patent, claim 26.

26. An integrated circuit device having at least one memory section which includes a plurality of memory cells, wherein the integrated circuit device outputs data on an external bus synchronously with respect to first and second external clock signals, the integrated circuit device comprises:

....

Delay locked loop circuitry to generate an internal clock signal using the first and second external clock signals

'804 patent, claim 26.

The claim language thus indicates that somehow the memory device is to use the information derived from the first and second clock signal to create an internal signal. One must consult the specification to understand how this is accomplished.

2. The Specification

In the "Background of the Invention" section, the specification relates that one "object of this invention is to provide a clocking scheme to permit high speed clock signals to be sent along the bus with minimal clock skew between devices." '918 patent, col. 3, ll. 27-29. "The two clocks together provide a synchronized high speed clock for all the devices on the bus." '918 patent, col. 8, ll. 29-30. Most significantly, in the "Clocking" subsection of the "Detailed Description," the inventors explain:

Clocking a high speed bus accurately without introducing error due to propagation delays can be implemented by having each device monitor two bus clock signals and then derive internally a device clock, the true system clock. The bus clock information can be sent on one or two lines to provide a mechanism for each bused device to generate an internal device clock with zero skew relative to all the other device clocks.

'918 patent, col. 18, l. 63 through col. 19, l. 4. This idea of clock skew can be best understood by reference to Figures 8a and 8b of the specification.

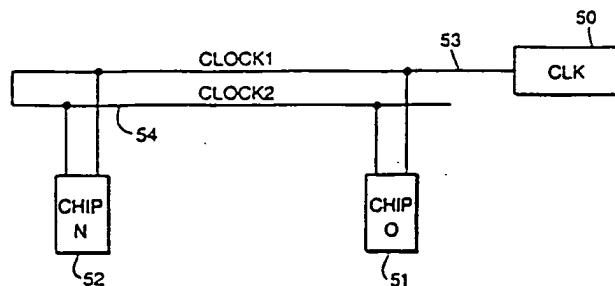


FIG. 8A

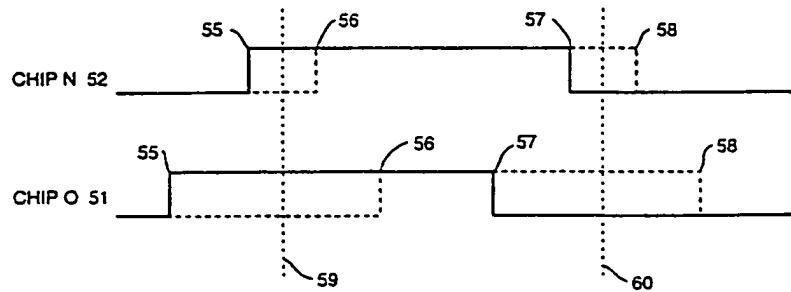


FIG. 8B

The specification clearly demonstrates how these figures represent the two clock signals:

Referring to FIG. 8a, in the preferred implementation, a bus clock generator 50 at one end of the bus propagates an early bus clock signal in one direction along the bus, for example on line 53 from right to left, to the far end of the bus. The same clock signal then is passed through the direct connection shown to a second line 54, and returns as a late bus clock signal along the bus from the far end to the origin, propagating from left to right. A single bus clock line can be used if it is left unterminated at the far end of the bus, allowing the early bus clock signal to reflect back along the same line as a late bus clock signal.

FIG. 8b illustrates how each device 51, 52 receives each of the two bus clock signals at a different time (because of propagation delay along the wires), with constant midpoint in time between the two bus clocks along the bus. At each device 51, 52, the rising edge 55 of Clock1 53 is followed by the rising edge of 56 of Clock2 54. Similarly, the falling edge 57 of Clock1 53 is followed by the falling edge 58 of Clock2 54. This waveform relationship is observed at all other devices along the bus. Devices which are closer to the clock generator have a greater separation between Clock1 and Clock2 relative to devices farther from the generator because of the longer time required for each clock pulse to traverse the bus and return along line 54, but the midpoint in time 59, 60 between corresponding rising or falling edges is fixed because, for any given device, the length of each clock line between the far end of the bus and that device is equal. Each device must sample the two bus clocks and generate its own internal device clock at the midpoint of the two.

'918 patent, col. 19, ll. 4-32. See also Fig. 13 of the '918 patent (showing how the rising and falling edges of the two bus clocks can be synchronized).

In essence, this portion of the specification explains how chips N and O, which are located in different positions along the bus lines, receive the clock signals at different points in time due to their locations relative to the origin of the clock signal. By reflecting the signal along a second line, the memory system can compensate for this delay and create a second clock signal. From these two signals, chips N and O create an internal clock signal which corrects the clock skew caused by propagation delay. In order to correct the skew, the two signals must necessarily contain different information, as Rambus' expert admitted. See Markman Hearing, Tr. pg. pp. 296-298 (testimony of Dr. Huber). Although the specification lists this clocking scheme as a preferred embodiment, it is actually the only embodiment of the clock in the entire specification. As with the analysis of the term "bus," it is significant that the specification limits the clock to a single embodiment. See generally Wang, 197 F.3d at 1380; Toro, 199 F.3d at 1301; O.I. Corp., 115 F.3d at 1581.

3. The Extrinsic Evidence

The constructions taught by the specification are confirmed by the testimony of Mr. McAlexander who explained that to one ordinarily skilled in the art that the Rambus clock scheme allows

the memory devices to sample each clock signal as it is received over the line and then averages the two signals such that every device is operating off the same clock, regardless of that device's location relative to the origin of the clock signal. Markman Hearing, Tr. pg. 457, line 2 to pg. 458, line 1 (testimony of Mr. McAlexander). Thus "the timing information and the difference between them is essential to this inventive concept of the clock design." Id.³⁸

³⁸ Mr. McAlexander's testimony also explains how the clocking scheme comes full circle to the primary objective of the invention and the use of a multiplexed bus:

In the prior art where the address information goes down one bus and data is responded to the bus on a totally different bus - so you have a data bus that's separate and distinct from the address bus - you could send down the control [or] the address information to a chip, activate it, . . . it goes in, finds the data from the storage cells and immediately sends it out to the data bus which is a separate bus.

In the multiplexed design, the data must share the same bus as the address information. And so the . . . controlling system must assure that at no time does address or control information reside on the bus at the same time that data . . . is coming back on the bus; otherwise, you would end up with a collision.

So in order to arbitrate that and to make sure that nothing is on the bus when it's not supposed to be, the whole system has to be in sync. Every system, every . . . chip, every component on the bus has to be operating under the exact same timing constraints.

That's why it's important and valuable . . . to use a clock design that will synchronize everything together.

In the interpretation of these terms, Rambus once again eschews the language of the specification, choosing to rely instead on the testimony of its expert who says that one skilled in the art would recognize that the first and second external clock signal can have, but does not need to have, different timing information in each signal. Therefore, according to Rambus, it is unnecessary to tie the claim definition to the language of the specification.³⁹ This approach runs afoul of the principle that the patent specification must always be reviewed to see if the patentee used the terms in a manner other than their ordinary meaning. Vitronics, 90 F.3d at 1582. Thus, even if one accepted Rambus' contention that the ordinary and accustomed meaning of first and external clock signals would be known to one of skill in the trade, the patent specification only describes a clocking scheme which corrects clock skew by creating an internal clock based on differing external clock signals.

4. Construction

Based on the claim language and the specification the term "first internal clock signal" is construed to mean "a periodic signal received by the memory device from an external source to

Markman Hearing, Tr. pg. 465 l. 6 to pg. 466, l. 10.

³⁹ See Testimony of Dr. Huber, pg. 302, l. 25 - pg. 303 l. 2 ("I don't need to go to the patent to interpret the term. Clock signal is a well known term").

provide first timing information." The term "second internal signal" is construed to mean "a periodic signal received by the memory device from an external source to provide second timing information that is different from the first timing information."

E. Integrated Circuit Device

Lastly, the parties contest the meaning of "integrated circuit device" as that term is found in claim 26 of the '804 patent. Rambus contends that the term means a "circuit constructed on a single monolithic substrate, commonly called a 'chip.'" Infineon, however, argues that representations made in the prosecution of the '804 patent limit this term to "a device composed of integrated circuits that include at least an ID register and related interface and comparison circuitry."⁴⁰

1. The Claim Language And The Specification

Neither the claim language nor the specification inform the present inquiry. Indeed, the specification mentions an integrated circuit only once.

⁴⁰ The parties have agreed that "integrated circuit device" is only disputed as it occurs in claim 26 of the '804 patent and not as it appears in the other patents. This is because the relevant file history limits the representations made to the PTO to only the '804 patent.

2. The File History

The file history of this claim is the only relevant category of intrinsic evidence. During the prosecution of Claim 26 in the '804 patent, (which was at that time, U.S. Patent App. 08/798,525 or the '525 application) Rambus expressly limited its claims by adding certain restrictions in order to overcome the PTO's prior art rejections. In response to the rejections, Rambus submitted new claims -- including the claim that ultimately issued as claim 26 of the '804 patent. Rambus argued to the PTO that the newly submitted claims were different from prior art because they all contained a device ID register and relevant interface and comparison circuitry limitations:

The new claims submitted in this Supplemental Preliminary Amendment have been added to more definitely and fully protect Applicants' invention. These newly submitted claims are directed to a memory device (or an integrated circuit having memory) having (1) an internal register for storing an identification value, (2) interface circuitry to receive a request on an external bus, and (3) comparison circuitry to determine whether the identification information in the request corresponds to the identification value in the internal register - wherein when the identification information corresponds to the identification value, the memory device responds to the request.

Supplemental Preliminary Amendment, U.S. Patent App. 08/798,525, pp. 11-12. The '804 patent issued subsequently. Thus, it appears that Rambus believed that its claims did not cover devices without a device ID register and relevant interface and comparison

circuitry. To allow Rambus to broaden its claim in the face of this restriction would defeat the public notice function of the patent history. In Hockerson-Halberstadt, Inc. v. Avia Group Intn'l, Inc., 222 F.3d 951 (Fed. Cir. 2000), the Federal Circuit explained:

[The inventor's] argument therefore reduces to a request for a mulligan that would erase from the prosecution history the inventor's disavowal of a particular aspect of a claim term's meaning. Such an argument is inimical to the public notice function provided by the prosecution history. The prosecution history constitutes a public record of the patentee's representations concerning the scope and meaning of the claims, and competitors are entitled to rely on those representations when ascertaining the degree of lawful conduct, such as designing around the claimed invention. . . . Were we to accept [the inventor's] position, we would undercut the public's reliance on a statement that was in the public record and upon which reasonable competitors formed their business strategies.

Id. at 957 (internal citations omitted).

"Absent qualifying language in the remarks, arguments made to obtain the allowance of one claim are relevant to interpreting other claims in the same patent." Digital Biometrics, 149 F.3d at 1347. Rambus claims to have presented such "qualifying language" in a footnote of the above-quoted representation to the PTO, which mentions two of the inventive technologies claimed in this suit:

The memory devices or integrated circuits having memory of the present invention may include additional and/or other inventive aspects, including, for example, delay lock loop circuitry and/or an internal register to

store a value which is representative of a number of clock cycles to transpire before the memory device responds to a read request. This "latency" register may be employed to control the timing of the output data after receipt of, for example, a read request. However it is noted that, in light of the July 27, 1998 Office Action and the rejection based on Weymouth, these additional and/or other inventive aspects, although forming a basis of patentability in their own right, will not be the focus of these Remarks.

Supplemental Preliminary Amendment, Patent App. No. 08/798,525, p. 12, n. 1. Rambus maintains that the express exclusion of the delay locked loop system and the latency invention from the scope of the attorney's remarks makes it "preposterous" to read a device ID register limitation into the claims currently before the court. Instead, during the Markman hearing in this case, Rambus' expert Dr. Huber contended that the limitation applies to every other claim of the patent (claims 1-25) but not claim 26. This conclusion, according to Dr. Huber, is an obvious conclusion based on the fact that claim 26 includes the inventive technologies mentioned as additional and/or inventive features in the footnote.

Notwithstanding Rambus' current attempt to carefully craft its limitations without much support in the patent history, the footnote does not imply that the statement excludes claim 26, but rather establishes that, in addition to the device ID register, Rambus believed that it claims possessed other inventive features. The last sentence of the footnote shows that Rambus chose not to rely on those additional inventive features when distinguishing the

claims from prior art. See id. ("these additional and/or other inventive aspects, although forming a basis for patentability in their own rights, will not be the focus of these Remarks"). Therefore, it is appropriate to read "integrated circuit device" as containing a device ID register, interface circuitry to receive a request from an external bus, and comparison circuitry to determine whether the identification information in the request corresponds to the identification register of the device.

Moreover, Rambus' suggested requirement that the integrated device be constructed on a single monolithic substrate is not supported by the specification and actually is undermined by the doctrine of claim differentiation. Claim 182 of the Preliminary Amendment to Patent App. No. 08/222,646, claim 6 of Patent No. 5,638,334 and claim 18 of Patent No. 5,657,481 all included the limitation that the device was "on a single semiconductor substrate." Claim 26 of the 804 patent contains no such limiting language and the doctrine of claim limitation warns against reading such a limitation into the disputed claim language unless the intrinsic evidence counsels otherwise.

3. Construction

Thus, the patent history supports the construction that an integrated circuit device, as used on claim 26 of the '804 patent, must have a device ID register, interface circuitry and comparison circuitry.

III. The Extrinsic Evidence Generally: The Experts

The claim construction here has been accomplished largely without resort to extrinsic expert evidence, notwithstanding that the parties presented expert testimony addressing each disputed term. Having reviewed that testimony, the Court found it useful mostly in understanding the meaning of technical terminology other than the disputed terms as that terminology is used in the claims and specification. See Pitney Bowes, 182 F.3d at 1309 ("it is entirely appropriate, perhaps even preferable, for a court to consult trustworthy extrinsic evidence to ensure that claim construction it is tending to from the patent file is not inconsistent with clearly expressed, plainly apposite, and widely held understanding in the pertinent technical field").

As outlined in the substantive discussion of each term construed, Rambus pressed constructions that generally found little, if any, support in the claim language or the specification, depending in significant part upon the expert testimony of Dr. Huber whose testimony was generally at odds with the statements made by the inventors in the claims and specification. Thus, his extrinsic evidence had to be substantially disregarded as contradictory of the intrinsic evidence. Also, it was difficult to credit Dr. Huber's testimony on disputed terms because it reflected the general, and disturbing, tendency of Rambus to distance its current constructions from what the inventors said in making the

claims and explaining the inventions in the specification,⁴¹ and, in so doing, to use the claim construction process to broaden claims, rather clearly not made in the intrinsic evidence.

The record here, and the approach to claim construction taken by Rambus, illustrate the wisdom and importance of the rules of law that establish a hierachal distinction between intrinsic and extrinsic evidence. On the other hand, the testimony of Mr. Joseph McAlexander, the expert offered by Infineon, was quite helpful and very credible because it was tethered closely to the intrinsic evidence and was not contradictory of the claim language or the specification. Notwithstanding that his testimony was reliable and informative it was ultimately not essential except as specifically cited in the construction. See Pitney Bowes, 182 F.3d at 1309 ("Although the patent file may often be sufficient to permit the judge to interpret the technical aspects of the patent properly, consultation of extrinsic evidence is particularly appropriate to ensure that his or her understanding of the technical aspects of the patent is not entirely at variance with the understanding of one skilled in the art").

⁴¹ Moreover, Dr. Huber left the impression that he was more an advocate than he was one generally knowledgeable in the field of the invention, notwithstanding his rather impressive curriculum vitae.

CONCLUSION

For the foregoing reasons, the disputed terms in the four patents in suit are to be construed as reflected herein.

The Clerk is directed to send a copy of this Memorandum Opinion to all Counsel of Record.

It is so ORDERED.

Robert E. Payne
United States District Judge

Richmond, VA

Date: *March 15, 2001*

07(1)4-42

To the
European Patent Office

Tabulation marks

I. Patent opposed		for EPO use only	
		Opp. No.	OPPO (1)
Patent No.		EP 1 004 956 B1	
		00101832.4	
Application No.		03.01.2001	
		Date of mention of the grant in the European Patent Bulletin (Art. 97(4), 99(1) EPC)	
Title of the invention: METHOD OF OPERATING A SYNCHRONOUS MEMORY HAVING A VARIABLE DATA OUTPUT LENGTH			
II. Proprietor of the Patent		RAMBUS INC.	
first named in the patent specification			
Opponent's or representative's reference (max. 15 spaces)		CST/M70615G(D1)	
III. Opponent		OREF	
Name		OPPO (2)	
Address			
MICRON EUROPE Ltd Micron House Wellington Business Park Dukes Ride Crowthorne Berkshire RG45 6LS UNITED KINGDOM			
State of residence or of principal place of business		UNITED KINGDOM	
Telephone/Telex/Fax			
Multiple opponents		<input checked="" type="checkbox"/> further opponents see additional sheet	
IV. Authorisation			
1. Representative (Name only one representative to whom notification is to be made)		OPPO (3)	
Name		Christopher Stephen Tunstall	
Address of place of business		Harrison Goddard Foote Tower House Merrion Way Leeds LS2 8PA UNITED KINGDOM	
Telephone/Telex/Fax		+44 113 290 1400	+44 113 244 2829
Additional representative(s)		<input type="checkbox"/> (on additional sheet/see authorisation)	
2. Employee(s) of the opponent authorised for these opposition proceedings under Art. 133(3) EPC		OPPO (5)	
Name(s): <i>Zur Kasse (A) € 1226,-</i>			
Authorisation(s)		<input checked="" type="checkbox"/> not considered necessary	
To 1./2.		<input type="checkbox"/> has/have been registered under No. _____	
		<input type="checkbox"/> is/are enclosed	



Notice of Opposition to a European Patent

To the
European Patent Office

Tabulation marks

for EPO use only

I. Patent opposed

Opp. No.	OPPO (1)
EP 1 004956 B1	

Patent No.

Application No.

Data of mention of the grant in the European Patent Bulletin
(Art. 97(4), 99(1) EPC)

Title of the invention:

II. Proprietor of the Patent

first named in the patent specification

Opponent's or representative's reference (max 15 spaces)

OREF

III. Opponent

OPPO (2)

Name

MICRON TECHNOLOGY ITALIA, S.R.L.

Address

Via Antonia Pacinotti 5/7
Nucico Industrial (AQ)
Building #2
87051 Avezzano (AQ) Italia
ITALYState of residence or of principal
place of business

ITALY

Telephone/Telex/Fax

Multiple opponents

 further opponents see additional sheet

IV. Authorisation

OPPO (9)

1. Representative
(Name only one representative to
whom notification is to be made)

Name

Address of place of business

Telephone/Telex/Fax

Additional representative(s)

 (on additional sheet/see authorisation)

OPPO (5)

2. Employee(s) of the opponent
authorised for these opposition
proceedings under Art. 133(3)
EPC

Name(s):

Authorisation(s)

To 1./2.

 not considered necessary has/have been registered
under No. _____ is/are enclosed

for EPO use only

V. Opposition is filed against

- the patent as a whole
- claim(s) No(s)



Claims 1-21

VI. Grounds for opposition:**Opposition is based on the following grounds:**

- (a) the subject-matter of the European patent opposed is not patentable (Art. 100(a) EPC) because.

- it is not new (Art. 52(1); 54 EPC)
- it does not involve an inventive step (Art. 52(1); 56 EPC)
- patentability is excluded on other grounds, i. e.

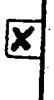
Art.



- (b) the patent opposed does not disclose the invention in a manner sufficiently clear and complete for it to be carried out by a person skilled in the art (Art. 100(b) EPC; see Art. 83 EPC).



- (c) the subject-matter of the patent opposed extends beyond the content of the application/ of the earlier application as filed (Art. 100(c) EPC; see Art. 123(2) EPC).

**VII. Facts and arguments**

(Rule 55(c) EPC)

presented in support of the opposition are submitted herewith on a separate sheet (annex 1)

**VIII. Other requests:**

That the Patent be revoked in its entirety.

That this opposition be subject to EXPEDITED HANDLING.

IX. Evidence presented

for EPO use only

Enclosed = will be filed at a later date =

A. Publications:

Publication data

- 1 International Patent Application WO91/16680, Published October 31, 1991 ("The Parent Application")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 2 US Patent 4,763,249, Published August 9, 1988 ("Bomba")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 3 US Patent 4,394,753, Published July 19, 1983 ("Penzel")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 4 US Patent 4,785,428, Published November 15, 1988 ("Bajwa")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 5 "The Scalable Coherent Interface Project (SuperBus)", SCI-22Aug88-doc ("SCI A")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 6 "Scalable Coherent Interface", SCI-28Nov88-doc20 ("SCI B")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 7 P1596: "SCI, A Scalable Coherent Interface", SCI-28Nov88-doc 2 ("SCI C")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

Continued on additional sheet

B. Other evidence

Continued on additional sheet

IX. Evidence presented

for EPO use only

Enclosed = will be filed at a later date =

A Publications:

Publication date

1 "Proposal for Clock Distribution in SCI" - 5/5/89 ("SCI D")

Particular relevance (page, column, line, fig.).

See Statement of Grounds of Opposition

2 Norsk Data Report - "A Proposal for SCI Operation" by Knut Aines - November 1988 ("SCI E")

Particular relevance (page, column, line, fig.).

See Statement of Grounds of Opposition

3 "Scalable I/O Architecture for Buses" by David V. James, SCI-28Nov88-doc3 (SCI F")

Particular relevance (page, column, line, fig.).

See Statement of Grounds of Opposition

4 Motorola MC88200 Cache/Memory Management Unit User's Manual, Published 1988 ("MC88200")

Particular relevance (page, column, line, fig.).

See Statement of Grounds of Opposition

5 Japanese Patent Application No. S63-142445, Published June 14, 1988, and English Translation ("Taguchi")

Particular relevance (page, column, line, fig.).

See Statement of Grounds of Opposition

6 US Patent 4,315,308, Published February 9, 1982 ("Jackson")

Particular relevance (page, column, line, fig.).

See Statement of Grounds of Opposition

7 Japanese Patent Application Sho 62-71428, Published October 5, 1988, and English Translation ("Yamaguchi")

Particular relevance (page, column, line, fig.).

See Statement of Grounds of Opposition

Continued on additional sheet

B. Other evidence

Continued on additional sheet

03-01-2001

EP001018324

OPPO

IX. Evidence presented

Enclosed = will be filed at a later date =

for EPO use only

A Publications:

Publication date

- 1 GigaBit Logic, 1989 GaAs IC Data Book & Designer's Guide, August 1989, 12G014
256'4-Bit Registered Self-Timed SRAM ("GigaBit")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 2 US Patent No. 4,499,536, Issued February 12, 1985 ("Gemma")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 3 Japanese Patent Application Sho 62-185253, Published January 31, 1989, and English Translation ("Kumagai")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 4 UK Patent Application GB-2,197,553, Published May 18, 1988 ("Lofgren")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 5 IEEE Journal of Solid State Circuits, Vol. 25, No. 1, February 1990, "An On-Chip Smart Memory for a Data-Flow CPU" ("Uvieghara")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 6 US Patent No. 4,637,018, Issued January 13, 1987, ("Flora"), as Exemplifying Common General Knowledge

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 7 Japanese Patent Application JP-A-01-284132, published November 15, 1989, and English Translation ("Kosugi")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

Continued on additional sheet

B. Other evidence

Continued on additional sheet

IX. Evidence presented

for EPO use only

Enclosed = will be filed at a later date =

A. Publications:

Publication date

- 1 Motorola MC88200 Cache/Memory Management Unit User's Manual, Published 1988, as Exemplifying Common General Knowledge ("MC88200")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 2 US Patent 4,680,738, Published July 14, 1987 ("Tam")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

- 3 US Patent 4,330,852, Published May 18, 1982 ("Redwine")

Particular relevance (page, column, line, fig.):

See Statement of Grounds of Opposition

4

Particular relevance (page, column, line, fig.):

5

Particular relevance (page, column, line, fig.):

6

Particular relevance (page, column, line, fig.):

7

Particular relevance (page, column, line, fig.):

Continued on additional sheet

B. Other evidence

Continued on additional sheet

X. Payment of the opposition fee is made

as indicated in the enclosed voucher for payment of fees and costs (EPO Form 1010)

for EPO use only

XI. List of documents

Enclosure
No.

No. of copies

- | | | |
|----|--|---|
| 0 | <input checked="" type="checkbox"/> Form for notice of opposition | <input type="text" value="2"/> (min. 2) |
| 1 | <input checked="" type="checkbox"/> Facts and arguments (see VII.) | <input type="text" value="2"/> (min. 2) |
| 2 | Copies of documents presented as evidence (see IX.) | <input type="text"/> |
| 2a | — Publications | <input type="text"/> (min. 2 of each) |
| 2b | — Other documents | <input type="text"/> (min. 2 of each) |
| 3 | <input type="checkbox"/> Signed authorisation(s) (see IV.) | <input type="text"/> |
| 4 | <input checked="" type="checkbox"/> Voucher for payment of fees and costs (see X.) | <input type="text" value="1"/> |
| 5 | <input type="checkbox"/> Cheque | <input type="text"/> |
| 6 | <input checked="" type="checkbox"/> Additional sheet(s) | <input type="text" value="2"/> (min. 2 of each) |
| 7 | <input type="checkbox"/> Other (please specify here): | <input type="text"/> |

XII. Signature
of opponent or representative

Place Leeds, UNITED KINGDOM

Date January 3, 2001



Christopher Stephen Tunstall

EUROPEAN PATENT NO. 1 004 956

STATEMENT OF GROUNDS OF OPPOSITION

1. ADDED SUBJECT-MATTER ART. 76(1) EPC

1.1. The Patent discloses subject-matter not disclosed in the Parent Application as originally filed on two counts, namely impermissible claim broadening and impermissible intermediate claim generalisations.

1.2. Impermissible Claim Broadening

1.2.1. The Parent Application disclosed and claimed several different alleged inventions. It is necessary to determine to which of these the granted claims relate, otherwise the comparison required by Art. 76(1) EPC cannot be made. In total, there were 7 objectives of the invention and 21 independent claims. The interrelationships between the objectives and claims need not be explored here in full. For present purposes, it is clear that granted claim 1 is based upon PCT claim 38, this being the broadest claim to modifiable data block size transfers. This was acknowledged by the patentee in its representative's letter dated 28 January 2000, numbered para. 1. It is also clear that neither PCT claim 38 nor granted claim 1 provides any solution to the 6 original objectives that have not found their way into the Patent.

1.2.2. The one objective of the alleged invention as set out in the Parent Application that survived in the Patent was this.

1.2.2.1. "One object of the present invention is to use a new bus interface built into semiconductor devices to support high-speed access to large blocks of data from a single memory device by an external user of the data, such as a microprocessor, in an efficient and cost-effective manner" [Parent Application, 6:8-12]. This will be referred to as "objective 1."

1.2.3. However, one of the 6 abandoned objectives, or rather one part of a compound objective, is relevant to the present discussion as follows.

1.2.3.1. "Yet another objective of this invention is to provide a method for transmitting address, data and control information over a relatively narrow bus ..." [Parent Application, 6: 21-23]. This will be referred to as "objective 2."

1.2.4. PCT claim 38 is a dependent claim. It depends from PCT claim 28. PCT claim 28 depends from PCT claim 26. PCT claim 26 depends from PCT claim 25. PCT claim 25 is an independent claim. In these circumstances, it is not altogether surprising that PCT claim 38 can be shown not to achieve any of the 7 primary objects of the invention. That is a privilege normally reserved for independent claims rather than claims three steps lower in the hierarchy.

1.2.5. If any objective was achieved by PCT claim 25, it was objective 2

1.2.5.1. PCT claim 25 reads as follows.

"25. A bus subsystem comprising:
two semiconductor memory devices connected in parallel to a bus,
wherein one of said semiconductor devices is a master device,
said master device including a means for initiating bus transactions,
said bus including a plurality of bus lines for carrying substantially
all address, data and control information needed by said devices,

said control information including device-select information,
said bus containing substantially fewer lines than the number of bits
in a single address, and

said bus carrying device-select information without the need for
separate device-select lines connected directly to individual devices on said
bus, whereby said master device initiates bus transactions which transfer
information between said semiconductor devices on said bus."

1.2.5.2.

This claim clearly relates to a bus subsystem. The bus includes a plurality of bus lines for carrying substantially all address, data and control information needed by the device for communication with substantially every other device connected to the bus, and has substantially fewer bus lines than the number of bits in a single address. A bus including a plurality of such general purpose lines, each carrying in a time-multiplexed manner substantially all address, data and control information needed by the device for communication with substantially every other device connected to the bus, and having substantially fewer bus lines than the number of bits in a single address, will be referred to in the present document as a "highly multiplexed bus." That the Parent Application was concerned with such a bus is constantly reinforced throughout the Parent Application including the summary of the invention and the beginning of the specific description [Parent Application, 7:10-19; 7:25-8:2; 11:16-25]. In addition, the bus subsystem of PCT claim 25 requires device selection to be accomplished using the control information carried by the bus.

1.2.5.3.

The bus is clearly a "relatively narrow bus" as required by objective 2. It is explicitly said to contain "substantially fewer bus lines than the number of bits in a single address." The operation of the bus subsystem of claim 25, with the master device initiating bus transactions over a bus that carries substantially all address, data and control information needed by the devices and contains substantially fewer lines than the number of bits in a single address, amounts to a "method for transmitting address, data and control information over a relatively narrow bus ..." as required by objective 2.

1.2.6. Objective 1 was not achieved by PCT claim 25

1.2.6.1.

Objective 1 concerns "high-speed access to large blocks of data from a single memory device ... in an efficient and cost-effective manner." PCT claim 25 makes no reference to the transfer of large blocks of data, nor indeed of blocks of data of any size. Nor does it make reference to memory devices. It contains no integers that in any way determine the ability of the bus subsystem to achieve objective 1. Whether or not the bus subsystem of claim 25 can meet objective 1 depends upon matters that do not form the subject of claim 25. These matters include the measures described in "DRAM Column Access Modification" [Parent Application, 59:4-62:2]. These measures quite clearly concern a "bus interface built into semiconductor devices." They include internal I/O multiplexing, allowing the interfacing of a memory device running at a relatively slow internal clock rate with the high-speed bus of the alleged invention [Parent Application, 60:1-6]. This increases the bandwidth of DRAM access [Parent Application, 59:19-2; 61:3-13]. It was acknowledged that the invention lay not in these measures per se, but in their use with the high-speed highly multiplexed bus of the invention [Parent Application 59:23-25]. These measures were the subject-matter of claims 82-90 and 114-123 of the Parent Application. This fits into the expected scheme of things: a number of these claims are independent.

1.2.6.2.

As has been shown, if any of the primary objectives of the invention was achieved by PCT claim 25, it was objective 2. It is hardly surprising that objective 1 was not achieved by this claim. One would not normally expect a

single independent claim to achieve two distinct and unrelated primary objectives of the invention.

1.2.7. Objective 1 was not achieved by PCT claim 26

1.2.7.1. PCT Claim 26 restricts PCT claim 25 to the case of a memory device connected to the bus, having one or more discrete memory sections and a modifiable address register to store memory address information which corresponds to the one, or each, discrete memory section. The bus is an integer of PCT claim 25. Whilst PCT claim 26 mentions a memory device, it says nothing that concerns "high-speed access to large blocks of data" from the memory device "in an efficient and cost-effective manner."

1.2.8. Objective 1 was not achieved by PCT claim 28

1.2.8.1. PCT claim 28 restricts PCT claim 26 to packet-based split-cycle transactions, in which the request packet from the bus master includes address and control information. The address information points to at least one memory location within a discrete memory section of the memory device. The memory section and memory device are integers of PCT claim 26. The control information includes information about the requested bus transaction and about the access time (corresponding to a number of bus cycles that need to intervene before bus access begins). PCT claim 28 says nothing that concerns "high-speed access to large blocks of data" from the memory device "in an efficient and cost-effective manner."

1.2.8.2. In relation to access times, an integer of PCT claim 28, there are few references in the Parent Application. One reference however reads as follows.

"A request packet and the corresponding bus access are separated by a selected number of bus cycles, allowing the bus to be used in the intervening bus cycles by the same or other masters for additional requests or brief bus accesses. Thus multiple, independent accesses are permitted, allowing maximum utilisation of the bus for transfer of short blocks of data. Transfers of long blocks of data use the bus efficiently even without overlap because the overhead due to bus address, control and access times is small compared to the total time to request and transfer the block" [Parent Application, 15:23-16:7].

1.2.8.3. This passage is a very clear statement of the advantages of modifiable access times. In the transfer of short blocks of data, modifiable access times allow for interleaving of requests and thus greater efficiency. In the transfer of long blocks of data, the need to transfer access time information on the bus is an overhead. It reduces efficiency. There is no possibility of interleaving long block requests. The preferred embodiments described in the Parent Application do not allow it. However, this overhead is an acceptable efficiency reduction because, compared with the overall length of the transaction, the overhead is small.

1.2.8.4. Objective 1 concerns "high-speed access to large blocks of data ... in an efficient and cost-effective manner." Modifiable access times contribute nothing to the achievement of that objective; as acknowledged in the Parent Application itself, they subtract from it. PCT claim 28 frustrates this objective.

1.2.9. Objective 1 was not achieved by PCT claim 38

1.2.9.1. PCT claim 38 restricts PCT claim 28 to situations in which the control information includes a block-size value that encodes and specifies the size of

the block of data to be transferred. The control information is an integer of PCT claim 28. Here at last, in a thrice dependent claim, an afterthought to an afterthought to an afterthought, is a reference to the transfer of a block of data. Claim 38 does not require that the size of the data block be large. The claim says nothing about whether the block is small or large. Moreover, claim 38 merely states that the size of the data block to be transferred is specified in the transaction request packet. It says nothing that determines whether the data block transfer is high-speed, efficient or cost effective, merely that its size can be determined. As discussed above, the answers to these questions lie in subject-matter that is entirely absent from the claim and is to be found in PCT claims 82-90 and 114-123.

- 1.2.10. To summarise, the granted claims find their basis in thrice dependent PCT claim 38. Independent PCT claim 25 from which PCT claim 38 ultimately depends, achieves (if anything) objective 2 and does not achieve objective 1. Objective 1 is not achieved by PCT claims 25 or 38 or any intervening claim, but is achieved by wholly different claims with wholly different subject-matter. PCT claim 38 and granted claim 1 achieve none of the primary objectives of the alleged invention.
- 1.2.11. The Technical Board of Appeal case law on claim broadening is settled.
- 1.2.11.1. The test for whether claim broadening is contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also, was clearly enunciated in decision T331/87, following decisions T194/84 and T260/85. The test was approved in decision T514/88 and is threefold [T331/87, Reasons:6].
- 1.2.11.2. "The removal of an integer from a claim may not violate Art. 123(2) EPC provided the skilled reader would directly and unambiguously recognise that:
1. the integer was not explained as essential in the original disclosure;
 2. it is not, as such, indispensable for the function of the invention in the light of the technical problem it serves to solve; and
 3. the replacement or removal requires no real modification of other integers to compensate for the change."
- 1.2.12. Granted claim 1 lacks an essential element of the alleged invention as originally disclosed.
- 1.2.12.1. Granted claim 1 relates to a method of operating a semiconductor memory device that includes, "... receiving block size information wherein the block size information defines an amount of data to be output onto an external bus in response to a read request; and outputting the amount of data ... in response to a read request ..." PCT claim 38, through its thrice dependent structure, explicitly required the bus to be a highly multiplexed bus. Granted claim 1, if it no longer contains this implied requirement, is broader than PCT claim 38. Such broadening is inadmissible.
- 1.2.12.2. Was the highly multiplexed bus explained as essential in the original disclosure?
- 1.2.12.2.1. According to decisions T260/85 and T527/88, in determining what is explained or disclosed in the application, it is necessary to read the document as a whole and individual passages within it in context. In view of this, a line-by-line analysis is no substitute for gaining an overall impression of the document by reading it from start to finish. However, the following points are mentioned to give a flavour for the disclosure of the Parent Application.

- 1.2.12.2.1.1. The Parent Application was entitled "Integrated Circuit I/O Using a High Performance Bus Interface".
- 1.2.12.2.1.2. The discussion of the prior art drew many distinctions between the numerous documents mentioned and the alleged invention, some of which were applicable to different alleged inventions from those originally claimed. However, in every case bar one, the documents were distinguished from the invention on the basis that it did not possess features of the bus interface [Parent Application, 3:14-18; 3:25-4:3; 4:5-7; 4:13-14; 4:16-18; 5:13-14; 6:5-7]. The single exception was a reference that described a clocking scheme that was distinguished from the clocking scheme used in the invention [Parent Application, 5:19-25]. The claims in the Parent Application that were directed to the clocking scheme, claims 73-81, and those directed to the chip package, claims 91-94, were the only claims not limited to the highly multiplexed bus.
- 1.2.12.2.1.3. Moreover, the state of the art was summed up thus: "None of the buses described in patents or other literature use only bused connections. All contain some point-to-point connections on the backplane" [Parent Application 5:13-15].
- 1.2.12.2.1.4. The summary of the invention opens with a discussion of the highly multiplexed bus [Parent Application, 7:10-19]. The necessary modifications that conventional DRAMs must undergo to comply with the alleged invention is described [Parent Application, 8:9-9:7] These modifications concern the interface with the highly multiplexed bus.
- 1.2.12.2.1.5. The objective achieved by PCT claim 38 via its dependency upon PCT claim 25, was concerned with the highly multiplexed bus. This is highly relevant [T514/88].
- 1.2.12.2.1.6. The entire specific description is concerned with a highly multiplexed bus [cf. Parent Application, 11:16-12:10].
- 1.2.12.2.1.7. All of the original claims, apart from those directed to the clocking and packaging schemes, were limited to the highly multiplexed bus.
- 1.2.12.2.2. All of these matters point to the essentiality of the highly multiplexed bus. Indeed, the whole tenor of the Parent Application was that it concerned a new bus architecture. There is no disclosure, suggestion or implication that anything else was contemplated. It is what the inventors contemplated, as objectively determined from their original application, that counts [T260/85, Reasons:10]. The Parent Application clearly explained that the highly multiplexed bus was an essential element of the alleged invention of PCT claim 38. No other reading or explanation is possible.
- 1.2.12.3. Was the highly multiplexed bus, as such, indispensable for the function of the invention in the light of the technical problem it serves to solve?
- 1.2.12.3.1. This is a straightforward question to answer. The technical problem that PCT claim 38 served to solve, via its dependency upon PCT claim 25, was to provide a method for transmitting address, data and control information over a relatively narrow bus ..." [Parent Application, 6: 21-23]. The only relatively narrow bus disclosed is the highly multiplexed bus. In this context, it is noteworthy that the acknowledged prior art

included buses in which the address information was multiplexed [Parent Application, 4:15-20]. It also included buses in which data and address information was multiplexed on the same lines [Parent Application, 4:21-23]. In each case, some control signals were bused. In distinguishing the alleged invention from the prior art by requiring "address, control and data" information to be transmitted on a "relatively narrow" bus, objective 2 can only have meant the highly multiplexed bus disclosed. Therefore, compatibility with the highly multiplexed bus was not just indispensable; it was a prerequisite.

1.2.12.3.2. Moreover, not only was the objective to be achieved expressly stated to concern compatibility with the highly multiplexed bus, but also any performance advantages attributable to modifiable block size transactions are present only in the context of the highly multiplexed bus.

1.2.12.4. Would the removal of the requirement for the highly multiplexed bus require any real modification of other integers to compensate for the change?

1.2.12.4.1. Again, this is a straightforward question to answer. Removal of this requirement has far-reaching consequences.

1.2.12.4.2. Firstly, the bus interface circuits of the semiconductor devices would no longer need to allow it to demultiplex and decode relevant bus transactions. They require modification.

1.2.12.4.3. Secondly, the nature of the "request packet" of PCT claim 38, via its dependency on PCT claim 28, is defined by the highly multiplexed nature of the bus. Removal of this interdependence would be a radical departure from the teaching of the application as filed, in which requests had to be encoded and multiplexed onto the bus lines in the way shown in and described with reference to Fig. 4 of the Parent Application [Parent Application, 21:21-24:2]. Any removal of the requirement for the highly multiplexed bus would radically alter the meaning of the term "request" as between PCT claim 38 and granted claim 1.

1.2.12.4.4. Thirdly, as discussed, modifiable block size transactions are disclosed only in the context of the highly multiplexed bus and do not in the inventors' view have an existence independent of it. Any removal of the requirement for the highly multiplexed bus would also remove the justification for modifiable block size transactions as perceived by the inventors.

1.2.12.5. For all these reasons, if the requirement in PCT claim 38 for the highly multiplexed bus is absent from granted claim 1, its absence fails all three limbs of the test established by the Technical Board of Appeal. It would amount to inadmissible claim broadening contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

1.2.13. Granted claim 1 lacks further essential elements of the invention as originally disclosed.

1.2.13.1. Apart from the highly-multiplexed bus, granted claim 1 makes no explicit reference to a plethora of other integers of PCT claim 38. These include the following.

1.2.13.1.1. The reference in PCT claim 25 to the bus carrying device select information without the need for separate device-select lines is not explicitly recited.

- 1.2.13.1.2. The reference in PCT Claim 26 to the memory device having a modifiable address register to store memory address information corresponding to one or more discrete memory sections is not explicitly recited.
- 1.2.13.1.3. The reference in PCT claim 28 to packet-based split-cycle transactions is not explicitly recited.
- 1.2.13.1.4. The reference in PCT claim 28 to the transaction request packets including address and control information is not explicitly recited.
- 1.2.13.1.5. The reference in PCT claim 28 to the control information of the request packet including information about the access time (corresponding to a number of bus cycles to intervene before bus access begins) is not explicitly recited.
- 1.2.13.1.6. The reference in PCT claim 38 to the control information of the request packet including the block-size value is not explicitly recited.
- 1.2.13.2. If these requirements of PCT claim 38 are missing from granted claim 1, then these omissions are inadmissible. All the integers omitted were present in the embodiments described. They could each be discussed in isolation, but it is more informative to look at the overall picture. In effect, claim 25 of the Parent Application has been rewritten as a method claim (to avoid having to mention the bus master) and then relieved of its very essence, namely the requirements for the highly multiplexed bus and for the bus to carry device select information. Into this wholly emasculated claim have been introduced a series of isolated integers including all of the three method steps of granted claim 1, none of which relate to the objective originally achieved by PCT claim 25. Clearly, it is not permissible for the patentee to remove from claim 25 the very integers that allow it to achieve the objective that the inventors had in mind for it. It amounts to abandoning the claim and reconstructing, *ex post facto*, a new claim including a selection of integers isolated from the original disclosure, whether from the description or the claims, irrespective of their relevance to the original primary objectives of the invention. That is not allowed.
- 1.2.13.3. Even if such an approach were admissible in principle, which the decisions of the Technical Board of Appeal cited above clearly demonstrate not to be the case, granted claim 1 contains impermissible intermediate generalisations, as will be discussed in the following section.

1.3. Impermissible Intermediate Generalisation

1.3.1. The Technical Board of Appeal case law on intermediate generalisation is settled.

- 1.3.1.1. The test for whether intermediate generalisation is contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also, was clearly enunciated in decision T284/94, following decision T17/86.
- 1.3.1.2. "An amendment of a claim by the introduction of a technical feature taken in isolation from the description of a specific embodiment is not allowable under Art. 123(2) EPC if it is not clear beyond any doubt for a skilled reader from the application documents as filed that the subject-matter of the claim thus amended provides a complete solution to a technical problem unambiguously recognisable from the application."
- 1.3.1.3. A number of other passages from this decision shed useful light on how this test is to be applied, in particular the following statements.

- 1.3.1.3.1. Referring to T17/86, an isolated technical feature "may be introduced into a claim without contravening Art. 123(2) EPC, provided that it is "evident beyond doubt to a skilled person reading the description that this isolated technical feature on its own enables the object in view to be achieved" " [T284/94, Reasons point 2.1.3, para. 2].
- 1.3.1.3.2. "In following this decision, the object to be achieved by the subject-matter of the amended claim has to be established as well as whether the claims define all means necessary for achieving this object" [T284/94, Reasons point 2.1.4, para. 1].
- 1.3.1.3.3. "Because of the fact that features disclosed in the context of a specific embodiment and added to a claim may achieve in an unambiguously recognisable manner an object different from that present in the introductory part of a description, it should further be ascertained whether such a further object is disclosed and whether it is clear beyond doubt for a skilled person reading the application as filed that the added technical features on their own achieve this further object" [T284/94, Reasons point 2.1.4, para. 3]. This makes clear the important distinction between objects that are relevant to claim broadening, i.e. the original objectives of the invention set out as such, and objects relevant to intermediate generalisation, which can be derived from other parts of the application, but only to the extent they are disclosed.
- 1.3.2. Granted claim 1 includes a technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.
- 1.3.2.1. PCT claim 38 required the block-size value to be included in control information that, together with address information, is bundled into a single request. That is exactly what is described in the Parent Application with reference to Fig. 4 [Parent Application, 21:21-22:10, esp. 22:9-10; 27:23-30]. No advantage is ascribed in the Parent Application to the subject matter of PCT claim 38, as distinct from the subject matter of the claims from which it depends.
- 1.3.2.2. Granted claim 1 contains no explicit requirement for the block size value to be included with address information in a transaction request.
- 1.3.2.3. As stated above, no advantage is ascribed in the Parent Application to the subject matter of PCT claim 38, as distinct from the subject matter of the claims from which it depends. This is a crucial point. According to the test established by the Technical Board of Appeal for admissibility of intermediate generalisation, it can only be justified if there is an advantage – a solution to a technical problem – disclosed in the Parent Application that is completely solved by the intermediate generalisation. If there is no problem and no solution, no advantage disclosed, the generalisation must fail. Such is the case here.
- 1.3.2.4. In spite of the above, it is not difficult to understand the purpose of PCT claim 38 in the context of a system where high-speed access to blocks of data is desired, and in which all bus transaction requests are six bus cycles long. It would be foolish to require one bus transaction to establish a block size and a second bus transaction to make the block transfer request. This would only increase the block request overhead to which the Parent Application refers [Parent Application, 15:23-16:7]. This purpose, namely allowing block size selection with no additional overhead, was not disclosed as such, but if it had been, it would not assist the patentee: it is only achieved by integrating the block size value into the transaction request, which explains why the inventors took such pains to allow both very small block sizes (block size 1)

and very large block sizes (block size 1024) to be represented by binary encoding using only four bits [Parent Application, 28:1-11]. It also explains the use of the word "encodes" in addition to "specifies" in PCT claim 38. Unless the block size value is so integrated, bus overhead will increase.

1.3.2.5. For this reason, if granted claim 1 does not require that block size information be integrated into a transaction request, then the claim fails the test established by the Technical Board of Appeal. It amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

1.3.2.6. It is worthy of note that it is not enough in the present context to say that the advantage or purpose of PCT claim 38 is to allow block size selection; that is tautologous. It amounts to saying that the advantage of an integer is its own existence; the problem to be solved by an integer its own provision. On that basis, every intermediate generalisation could be justified; but that is not the law.

1.3.3. Granted claim 1 includes a further technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.

1.3.3.1. PCT claim 38, via its dependency on PCT claim 28, specifies, "... a means for said master device to request said memory device to prepare for a bus transaction by sending a request packet along said bus ..." PCT claim 28 and PCT claim 38 then go on to specify the information contained within the request packet. Clearly, PCT claims 28 and 38 define in general terms a bus transaction protocol. As such, it is applicable to all bus transactions. Accordingly, every bus transaction is initiated by the master device requesting the memory device to prepare for a bus transaction by sending a request packet along the bus, the request packet including the specified information. This applies as much to write transaction requests as it does to read transaction requests. This is consistent with the preferred embodiment of the invention, in which exactly the same form of transaction request, as shown in Fig. 4, is used to request read or write block transfers. The only difference is one bit identifying whether the transaction is a read or a write transaction [Parent Application, 22:19-23:3].

1.3.3.2. Notably, according to claim 28, and therefore claim 38 also, the request packet should contain access time information (corresponding to a number of bus cycles to intervene before bus access begins). As discussed, the technical problem solved by modifiable access times, as claimed in PCT claim 103, was to maximise bus utilisation for exchange of small blocks of data in a highly multiplexed system [Parent Application, 16:1-3]. Modifiable access time DRAM read operations were discussed in the Parent Application in conjunction with modifiable access time DRAM write operations. Whereas bus utilisation may be improved by providing for modifiable access times in respect of read operations only, it is not a complete solution to the problem of maximising bus utilisation. For the achievement of that objective, the device must have programmable access times in respect of all operations that require it to utilise the bus in response to a request. This underlines the fact that PCT claim 38 was for good reason not limited to a particular type of transaction request.

1.3.3.3. Granted claim 1 relates to a method in which a semiconductor memory device responds in a certain way to read requests only. If granted claim 1 were limited to the operation of read only devices such as ROMs, this limitation to read requests only would not be objectionable. However, granted claim 1 also covers the operation of read/write devices such as

RAMs that receive block-size values for use in read requests only, and not in write requests. This represents an impermissible intermediate generalisation.

1.3.3.4. As has already been mentioned, no technical problem that is disclosed in the Parent Application is solved by the subject matter of claim 38. According to the test established by the Technical Board of Appeal for admissibility of intermediate generalisation, it can only be justified if there is an advantage – a solution to a technical problem – disclosed in the Parent Application that is completely solved by the intermediate generalisation. If there is no problem and no solution, no advantage disclosed, the generalisation must fail. Such is the case here.

1.3.3.5. Even if the purpose behind PCT claim 38 as discussed above had been disclosed, the position would be no different. The purpose is to allow block size selection with no additional overhead. That applies just as much to write block size selection as to read block size selection.

1.3.3.6. For this reason, the introduction into granted claim 1 of the requirement for block size information in respect of read operations only fails the test established by the Technical Board of Appeal. It amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

1.3.4. Granted claim 3 includes a technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.

1.3.4.1. Granted claim 3 introduces into granted claim 1 the requirement for access time information that is found in PCT claim 28 or PCT claim 103. As discussed above, PCT claim 28 specifies information contained within a request packet. It defines in general terms a bus transaction protocol. As such, it is applicable to all bus transactions. Accordingly, every bus transaction is initiated by the master device requesting the memory device to prepare for a bus transaction by sending a request packet along the bus, the request packet including the specified information, in particular access time information, irrespective of whether it is a read or a write request. This is consistent with the preferred embodiment of the invention, in which exactly the same form of transaction request, as shown in Fig. 4, is used to request read or write block transfers and includes access time information [Parent Application, 27:1-15]. This type of read request access time information points to one of two access time registers that have been set up in advance.

1.3.4.2. On the other hand, according to PCT claim 103, "data may be transmitted to said [access time] register via said bus which establishes a predetermined amount of time that said semiconductor device thereafter must wait before using said bus in response to a request." This is more akin to the subject matter of granted claim 4, but as claim 4 depends from claim 3, it is relevant to claim 3 also. This type of access time information is what is put into the access time registers in the first place. It is clear from the language of PCT claim 103 that it is speaking of access time registers being set up prior to a request being made. This is consistent with the preferred embodiment of the invention, in which two access time registers are set up in advance. It is also clear from the language of PCT claim 103, in particular the phrase "thereafter must wait ... in response to a request," that it applies to all subsequent requests. This is also consistent with the preferred embodiment, in which transaction requests point to one of the two access time registers in all cases [Parent Application, 27:1-15].

1.3.4.3. This disclosure of access time information applying in all case to both read and write transactions can easily be understood. The advantage of a

modifiable access time register is, as discussed above, to allow "maximum utilisation of the bus for transfer of short blocks of data" [Parent Application, 15:1-3]. Maximum utilisation of the bus can only be achieved if all the devices connected to the bus use a programmable access time register to time their response to all requests directed to them.

- 1.3.4.4. Granted claim 3 requires the semiconductor memory device to respond in accordance with the access time information to read requests only. If granted claim 1 were limited to read only devices such as ROMs, this limitation to read requests only would not be objectionable. However, granted claim 1 also covers read/write devices such as RAMs that have programmable access times for read requests only, and not for write requests. This represents an impermissible intermediate generalisation.
- 1.3.4.5. As discussed above, the technical problem solved by modifiable access times, as claimed in PCT claims 28 and 103, was to maximise bus utilisation for exchange of small blocks of data in a highly multiplexed system [Parent Application, 16:1-3]. Modifiable access time DRAM read operations were discussed in the Parent Application in conjunction with modifiable access time DRAM write operations. Whereas bus utilisation may be improved by providing for modifiable access times in respect of read operations only, it is not a complete solution to the problem of maximising bus utilisation. For the achievement of that objective, the device must have programmable access times in respect of all operations that require it to utilise the bus in response to a request. Granted claim 3 is not so limited.
- 1.3.4.6. For this reason, the introduction into granted claim 3 of the requirement for modifiable access times in respect of read requests only fails the test established by the Technical Board of Appeal. It amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.
- 1.3.5. Granted claims 14 and 16 include a technical feature that has been isolated from its original context and does not provide a complete solution to any technical problem disclosed in the Parent Application.
- 1.3.5.1. Claims 14 and 16 of the patent require internal clock signals to be generated by "a delay locked loop". The term "delay locked loop" was never used in the Parent Application. It is a wholly new term. No delay locked loop was ever disclosed. The nearest thing to it was probably to be found somewhere in the circuit illustrated in Fig. 12 of the Parent Application. The broadest disclosure of that circuit in the Parent Application is to be found in claims 78 and 108, but these claims make no mention of anything that could be regarded as a delay locked loop. The broadest disclosure in the Parent Application of any sufficiently concrete functional detail of the circuit of Fig. 12 is to be found in claim 79.
- 1.3.5.2. Clearly, the object to be achieved by the circuit of Fig. 12 was to generate an internal clock signal synchronised to a time half way between the early and late bus clock signals [Parent Application, 46:20-47:1; 47:21-48:3]. This objective is only achieved if the integers of claim 79 of the Parent Application are present. Granted claims 14 and 16 are not so limited.
- 1.3.5.3. For this reason, the introduction into granted claims 14 and 16 of a delay locked loop in isolation fails the test established by the Technical Board of Appeal. It amounts to inadmissible intermediate generalisation contrary to Art. 123(2) EPC, and therefore Art. 76(1) EPC also.

2. ENTITLEMENT TO DIVISIONAL STATUS

- 2.1. For precisely the reasons discussed above and because the objectionable added subject matter was present in the Divisional Application as filed (a fundamental error), it follows that the Divisional Application and the Patent are not entitled to the benefit of the filing and priority dates of the Parent Application. They take as their filing and priority dates the date on which it was actually filed, namely January 29, 2000. No subsequent amendment can affect the position under Art. 76(1) EPC, as indicated in decision T873/94, Reasons, Section 1, Para. 4.

3. LACK OF NOVELTY ART. 54 EPC

- 3.1. Each and every claim of the Patent, having a priority date of January 29, 2000, lacks novelty.
- 3.2. WO91/16680 ("The Parent Application"), published October 31, 1991
- 3.2.1. The Parent Application discloses every feature disclosed in the Patent, including all the integers of every claim.
- 3.3. The matter disclosed in the Parent Application does not support any claim purporting to cover buses other than the highly multiplexed bus described or read requests other than read request packets. In proceedings brought in the UK, France and Germany under the patent arising from the Parent Application, the patentee has asserted that the terms "external bus" and "request" in granted claim 1 of that patent should be afforded an interpretation wider than that supported by the Parent Application. The following discussion assumes that such a wider interpretation may be advanced by the patentee in respect of the Patent.
- 3.4. The subject-matter of claims 1-5, 8-10, 12-15, and 17-21 of the Patent lacks novelty even if entitled to the declared priority date of April 18, 1990.
- 3.5. US Patent 4,763,249, Published August 9, 1988 ("Bomba")
- 3.5.1. Bomba discloses a system that includes a plurality of bus devices interconnected by a synchronous, multiplexed bus. The bus device can be constructed as a memory device with a plurality of storage locations and interconnection circuitry [Bomba, abstract]. The interconnection circuitry forms an integral part of the memory device [Bomba, 6:63-65]. Thus, Bomba discloses a "semiconductor memory device having at least one memory array which includes a plurality of memory cells" and its method of operation.
- 3.5.2. A master clock 144 connected to the communications path generates clocking signals (time and phase signals) for the bus devices [Bomba, 9:14-32 and Figs. 3A and 3B]. These signals are received by, *inter alia*, the interconnection circuitry of the memory device. Thus, Bomba discloses "receiving [in the memory device] an external clock signal having a fixed frequency."
- 3.5.3. A two-bit data length code is placed on the highest bits of data lines D[31:30] during a command/address cycle of a read transaction [Bomba, 13:59-14:4 and Fig. 4A]. The lower 30 bits contain the device "address" (the address is the 30 bit storage location where the transaction is to take place) [Bomba, 13:59-14:4 and 14:13-15]. The data length code specifies the length of the data transfer that is to take place, e.g. one to four cycles of 32 bit data [Bomba, 15:18-28]. The operation code for the read command is transmitted over information lines I[3:0] at the same time [Bomba, 13:59-61]. As can be seen from Figs. 1A, 2 and 3B of Bomba, each of these signal is received by, *inter alia*, the interconnection circuitry of the memory. Thus, Bomba discloses "receiving block size information, wherein the block size information defines an amount of data to be output onto an external bus in response to a read request."

- 3.5.4. A memory device on the bus, containing the relevant address, confirms receipt of the command/address cycle and may begin the data transaction [Bomba, 15:49-63]. The memory device outputs data on the data lines D[31:0] during a first data cycle. The memory device continues to output new data for as many data cycles as are specified during the command/address cycle [Bomba, 15:29-16:58]. Thus, Bomba discloses "outputting the amount of data corresponding to the block size information, in response to a read request." Data is generally placed on the data lines at the leading edge of internal clock TCLK [Bomba, 9:51-54]. TCLK is the transmitting clock generated locally in the memory device from the time and phase components of master clock 144 [Bomba, 8:46-63 and 9:27-39 and Fig. 3B]. Thus, the data is output "synchronously with respect to the external clock signal."
- 3.5.5. It follows that the subject-matter of claim 1 is not new.
- 3.5.6. The memory device outputs data on the data lines D[31:0] during a first data cycle. The memory device continues to output new data for as many data cycles as are specified during the command/address cycle [Bomba, 15:29-16:58]. The time between successive TCLK signals defines a cycle [Bomba, 9:40-41]. TCLK is generated locally in the memory device from the time and phase components of master clock 144 [Bomba, 8:46-63 and 9:27-39 and Fig. 3B]. Thus, Bomba discloses that "the amount of data corresponding to the block size information is output synchronously during a plurality of clock cycles of the external clock signal."
- 3.5.7. It follows that the subject-matter of claim 2 is not new.
- 3.5.8. Data is placed on the data lines at the leading edge of TCLK [Bomba, 9:51-54]. As stated above, TCLK is generated locally from the time and phase components of Master Clock 144. Fig. 3A of Bomba shows the correlation between TCLK and the external clock signals TIME (+) and PHASE (+). A rising edge of TCLK corresponds with a next rising edge of TIME (+) after a rising edge of PHASE (+). Thus, Bomba discloses that "data is output onto the external bus synchronously with respect to a rising edge transition of the external clock signal."
- 3.5.9. It follows that the subject-matter of claim 8 is not new.
- 3.5.10. Fig. 3A of Bomba also shows the correlation between TCLK and the external clock signals TIME (-) and PHASE (-). A rising edge of TCLK corresponds with a next falling edge of TIME (-) after a falling edge of PHASE (-). Thus, Bomba discloses that "data is output onto the external bus synchronously with respect to a falling edge transition of the external clock signal."
- 3.5.11. It follows that the subject-matter of claim 9 is not new.
- 3.5.12. Bomba discloses the use of two bits to represent the length of 4 different block sizes [Bomba, 15:18-24]. Thus, the two bits are a binary representation of the block size in units of 32 bits. Therefore, Bomba discloses that "the block size information is a binary representation of the amount of data to be output after receipt of a read request."
- 3.5.13. It follows that the subject-matter of claim 10 is not new.
- 3.5.14. As discussed above, a two-bit data length code is placed on the highest bits of data lines D[31:30] during a command/address cycle of a read transaction [Bomba, 13:59-14:4 and Fig. 4A]. The lower 30 bits contain the device "address" [Bomba, 13:59-14:4 and 14:13-15]. The operation code for the read command is transmitted over information lines I[3:0] at the same time [Bomba, 13:59-61].
- 3.5.15. It follows that the subject-matter of claim 12 is not new.

3.5.16. As stated above, the first internal clock signal *TCLK* is generated locally in each device on the bus using the master clock 144 [Bomba, 8:54-63]. Data is placed on the data lines at the leading edge of *TCLK* [Bomba, 9:51-54]. Thus, Bomba discloses "generating a first internal clock signal using the external clock signal wherein the amount of data corresponding to the block size information is output on to the external bus synchronously with respect to the first internal clock signal."

3.5.17. It follows that the subject-matter of claim 13 is not new.

3.5.18. Each memory device contains a large number of configuration registers (200-216). One such register is a control and status register 202 illustrated in Fig. 7C, that contains, *inter alia*, device ID information that is loaded into bits CSR[3:0] on system power-up or during a subsequent initialisation sequence [Bomba, 22:29-36]. Thus, Bomba discloses "a programmable identification register [in the memory device] to store an identification value to identify the memory device from a plurality of other memory devices on the external bus."

3.5.19. Device ID information is transmitted as part of a read transaction received by the memory device [Bomba, 8:23-24]. It forms a part of the thirty bit device address transmitted on lines D[29:0] [Bomba, 14:13-18]. Thus Bomba discloses "receiving [in the memory device] identification information."

3.5.20. Only the memory device addressed by the device address responds to the read transaction. The memory device compares its device ID with the ID information transmitted with the device address on lines D[29:0] and responds only if there is a match. Thus, Bomba discloses "determining whether the identification information corresponds to the identification value stored in the programmable identification register wherein, when the identification information corresponds to the identification value, the amount of data corresponding to the block size information is output onto the external bus synchronously with respect to the external clock signal."

3.5.21. It follows that the subject-matter of claim 17 is not new.

3.5.22. As stated above, the device ID information is loaded into bits CSR[3:0] on system power-up or during a subsequent initialisation sequence [Bomba, 22:29-36]. Thus Bomba discloses that "the programmable identification register is programmed" either "after power is applied to the memory device during an initialisation sequence" or "during an initialisation sequence of the memory device."

3.5.23. It follows that the subject-matter of claims 19 and 20 is not new.

3.6. US Patent 4,394,753, Published July 19, 1983 ("Penzel")

3.6.1. Penzel describes a highly integrated memory module. As illustrated in Fig. 1 of Penzel, the highly integrated memory module is a semiconductor memory device that includes a memory cell array (DECODER MEMORY) arranged as a plurality of rows and columns [Penzel, 2:45-53]. Thus, Penzel discloses a "semiconductor memory device having at least one memory array which includes a plurality of memory cells" and its method of operation.

3.6.2. As illustrated in Fig. 5 of Penzel, during a chained memory access (block read/write), the column address strobe signal *CAS* is pulsed [Penzel, 6:44-50]. Fig. 5 shows *CAS* as being a fixed frequency signal. Thus, Penzel discloses "receiving an external clock signal having a fixed frequency."

3.6.3. The memory device of Penzel includes a mode register [Penzel, 3:9-11]. The mode register is a nine bit register in which two bits *M*₀, *M*₁ determine whether the read is a $\times 1$ read, a $\times 4$ read or a $\times 8$ (or $\times 9$) read [Penzel, 3:37-42]. Two further bits, *M*₂, *M*₃

determine the number of chained accesses in a read operation (block size) [Penzel, 3:42-44 and 3:46-57]. The mode register is programmed externally via package pins P_0-P_8 [Penzel, 3:63-65]. Thus, Penzel discloses "receiving block size information, wherein the block size information defines an amount of data to be output onto an external bus in response to a read request."

- 3.6.4. The block read operation is described [Penzel, 4:21-26; 6:34-65]. As is clearly illustrated in Fig. 5, successive data is output in response to successive pulses of the external clock signal CAS. Thus, Penzel describes "outputting the amount of data corresponding to the block size information, in response to a read request, synchronously with respect to the external clock signal."
- 3.6.5. It follows that the subject-matter of claim 1 is not new.
- 3.6.6. As stated above, successive data is output in response to successive pulses of the external clock signal CAS. Thus, Penzel discloses that "the amount of data corresponding to the block size information is output synchronously during a plurality of clock cycles of the external clock signal."
- 3.6.7. It follows that the subject-matter of claim 2 is not new.
- 3.6.8. As shown in Fig. 5 of Penzel (understood with the necessary modifications in the case of a read operation), data is input or output in response to successive falling edges of CAS. Thus, Penzel discloses that "data is output onto the external bus synchronously with respect to a falling edge transition of the external clock signal."
- 3.6.9. It follows that the subject-matter of claim 9 is not new.
- 3.6.10. The bits M_2, M_1 of Penzel designate 1 [0,0], 2 [0,1], 4 [1,0] or 8 [1,1] accesses in a chain [Penzel, 3:47-57]. This is a logarithmic binary representation. Thus, Penzel discloses that "the block size information is a binary representation of the amount of data to be output after receipt of a read request."
- 3.6.11. It follows that the subject-matter of claim 10 is not new.
- 3.7. US Patent 4,785,428, Published November 15, 1988 ("Bajwa")
- 3.7.1. Fig. 1 of Bajwa is the functional block diagram of a DRAM controller (consisting of functional blocks 9, 11-13, 15-17) that controls a DRAM array 14. In the preferred embodiment, the DRAM controller itself is a chip, [Bajwa, 2:56], while the DRAM array 14 is a 1Mb conventional DRAM [Bajwa, 4:24]. The combination of the DRAM controller and the DRAM array constitutes a semiconductor memory device that interfaces with a synchronous bus 10. Thus, Bajwa discloses "a semiconductor memory device having at least one memory array which includes a plurality of memory cells" and its method of operation.
- 3.7.2. The DRAM controller performs the function of the bus interface of the memory device and is a clocked, self-timed device. The DRAM controller is driven by a clock having a two non-overlapping phase design [Bajwa, 2:52-54]. Four clocks PH1 and PH2 and their inverses PH1I and PH2I are disclosed as controlling the logic function [Bajwa, 2:54-59]. Fig. 4 shows how a memory access sequence is clocked with respect to PH1.
- 3.7.3. Bajwa further describes a clock management circuit 28 and RAS, WE, OE, and CAS pin control logic which receive an external clock in the form of clock signals CLKA, CLKB, CLKC and CLKD [Bajwa, Fig. 2]. The clock management circuit in combination with the clocking signals controls timing for the memory access [Bajwa, 6:18-23]. Thus, Bajwa discloses "receiving an external clock signal having a fixed frequency."

- 3.7.4. The DRAM controller provides timing sequences for one to four word memory accesses [Bajwa, 8:16-19]. During a memory operation, the DRAM sequence controller 24 receives a two bit block size information code (NUMWORDS) which corresponds to the number of words for the memory access (1 to 4 words) [Bajwa, 3:63-66; Fig. 2]. A 3-bit unary encoded or 2-bit binary encoded operation code signal is supplied to a SIGNAL MUX 22 to specify whether the operation is a read, write or refresh [Bajwa, 4:6-10]. Thus, Bajwa discloses "receiving block size information, wherein the block size information defines an amount of data to be output onto an external bus in response to a read request."
- 3.7.5. The program RAM 20, containing the DRAM access protocols and timings, jumps to the portion of program RAM memory specifying the access protocol and timing for the number of words specified in NUMWORDS [Bajwa, 8:14-38]. The signal multiplexer 22 selects the appropriate signals from the timing sequence stored in the program RAM 20 and drives the various RAS, CAS, OE, and WE signals, along with the signals from clock management unit 28, clocking signals (CLKA-CLKD) and ARREN [Bajwa, Figs. 1 and 2]. As shown these signals are synchronous with respect to, at least, clock PH1. They are supplied to the DRAM and cause the DRAM to output the selected amount of data. The read data coming in from the DRAM is queued in the memory control unit and is output onto the external bus (AP Bus) according to a pre-programmed timing sequence which indicates the clock cycle during which the transmission of the read data may begin [Bajwa, 5:64-6:14]. The AP bus timing requires the transmission of a data word on every bus cycle [Bajwa, 6:1-2]. Thus, the data is read out synchronously with respect to the external clock signal. Accordingly, Bajwa discloses "outputting the amount of data corresponding to the block size information, in response to a read request, synchronously with respect to the external clock signal."
- 3.7.6. It follows that the subject-matter of claim 1 is not new.
- 3.7.7. The read data from the DRAM is queued in the memory control unit and is output onto the external bus (AP Bus) according to a pre-programmed timing sequence which indicates the clock cycle during which the transmission of the read data may begin [Bajwa, 5:64-6:14]. The AP bus timing requires the transmission of a data word every bus cycle [Bajwa, 6:1-2]. Thus, when more than one data word is read from the DRAM, "the amount of data corresponding to the block size information is output synchronously during a plurality of clock cycles of the external clock signal."
- 3.7.8. It follows that the subject-matter of claim 2 is not new.
- 3.7.9. The memory includes a program RAM 20 for storing "timing sequences programmed into" the DRAM controller to match the speed of any number of existing arrays [Bajwa, 2:44-45; 2:6-7]. The timing sequences are programmed into program RAM 20, which is a 53x35 array as shown in Fig. 2. Each of the 53 rows is used for programming all 35 internal transition control signals required for one clock cycle. The 53 rows are used as follows. The DRAM controller is capable of block size access of 1 to 4 words in size [Bajwa, 8:28-34], and the timing sequence for each of the 4 block size types is programmed separately, i.e. by using:
- 3.7.9.1. 8 rows for the timing sequence of 1-word access, [Bajwa, Fig. 6] (max delay of 8 clock cycles),
 - 3.7.9.2. 11 rows for programming the timing sequence of a 2-word access (max delay of 11 clock cycles),
 - 3.7.9.3. 15 rows for programming the timing sequence of a 3-word access (max delay of 15 clock cycles),

- 3.7.9.4. 19 rows for programming the timing sequence of a 3-word access (max delay of 19 clock cycles),
- 3.7.10. This requires a total of 53 rows. Thus, program RAM 20 comprises at least four access-time registers used to specify 4 different access times corresponding to each of the 4 possible (programmable) block sizes. In this scheme, one pre-programmed signal RPYNOW signals the memory control unit to transmit the read reply packet to the AP bus [Bajwa, 5:64-6:14]. RPYNOW can correspond to any one of a number of clock cycles [Bajwa, 6:10-15]. Thus, Bajwa discloses "receiving access-time information wherein the access-time information is representative of number of cycles of the external clock signal to transpire before data is output onto the external bus after receipt of a read request." This information is programmed into a "programmable access time register."
- 3.7.11. It follows that the subject-matter of claims 3-5 is not new.
- 3.7.12. The block size information NUMWORDS is two bits [Bajwa, 3:63-66; Fig. 2]. These two bits correspond to block sizes of one to four words. Thus, NUMWORDS is a "binary representation of the amount of data to be output."
- 3.7.13. It follows that the subject-matter of claim 10 is not new.
- 3.7.14. When a read request is received, the central sequencing logic (CSL) 17 of the DRAM controller starts a timing sequence that provides all the timings required for accessing the DRAM array in each clock cycle. A SEQUENCE signal is set high to indicate that the DRAM controller is currently processing the request [Bajwa, 3:66-67; Fig. 4]. SEQUENCE is set low (inactive) when the processing of the request is completed. A RAS PRECHARGE is discussed [Bajwa, 7:15-26]. The DRAM RAS precharge time is specified by a 3-bit counter in the preferred embodiment. The counting of the precharge cycles is enabled when SEQUENCE goes inactive [Bajwa, 7:21-22]. Hence, the memory array is "automatically precharged after the read request has been executed" by the DRAM controller. When the count expires, precharge is complete and the DRAM controller is ready to process the next request (the precharge counter is also reloaded in preparation for the next read cycle) [Bajwa, 7:23-27].
- 3.7.15. It follows that the subject-matter of claim 21 is not new.
- 3.8. IEEE Standard for a Simple 32-Bit Backplane Bus: NuBus - ANSI/IEEE Std 1196-1987 ("NuBus")
- 3.8.1. The NuBus standard is a synchronous computer backplane bus standard in which the bus is used to connect devices and to provide certain resources to the connected devices [NuBus, Fig. 1]. Conventional memory modules are one type of device that may be attached to the NuBus [NuBus, p. 42 (A.1 Note)]. Thus, NuBus discloses "a semiconductor memory device having at least one memory array which includes a plurality of memory cells" and its method of operation.
- 3.8.2. The NuBus modules receive an external clock source (central system clock) [NuBus, Fig. 1; p. 4 (Section 2.1.1)]. The clock source is common to each NuBus module. The clock is a fixed frequency clock, nominally operating at 10 MHz, and is used to synchronise bus arbitration and data transfers [NuBus, p. 4 (Section 2.1.1)]. The clock signal is driven from one end of the bus to termination at the other end [NuBus, p. 44 (Section A.7)]. Thus NuBus describes "receiving an external clock signal having a fixed frequency."
- 3.8.3. NuBus provides for block data transfers in block sizes of 2, 4, 8 and 16 words [NuBus, p. 11 (Section 3.1.4)]. The number of data words transferred is controlled by the master and communicated during the start cycle. The block size and block starting address is transmitted over the 32 bit multiplexed address and data lines

03-01-2001

EP00101832.4

18

OPPO

(which run between the master and slave devices) while the START* signal is asserted [NuBus, Fig. 4; pp. 11-13]. Thus, NuBus discloses "receiving block size information, wherein the block size information defines an amount of data to be output onto an external bus in response to a read request."

- 3.8.4. The slave device drives the first word of the requested data onto the 32 bit multiplexed address and data lines. The start of the data transfer is synchronous with respect to CLK* (occurring at the rising edge) [NuBus, Fig. 5; p. 12]. Blocks of data are output from the slave synchronously with respect to CLK* until the desired block size is reached [NuBus, Fig. 5]. Thus, NuBus discloses "outputting the amount of data corresponding to the block size information, in response to a read request, synchronously with respect to the external clock signal."
- 3.8.5. It follows that the subject-matter of claim 1 is not new.
- 3.8.6. As stated above, the slave device drives the subsequent words of the requested data onto the 32 bit multiplexed address and data lines. Blocks of data are output from the slave synchronously with respect to CLK* until the desired block size is reached [NuBus, Fig. 5; p. 12]. Thus, NuBus discloses that "the amount of data corresponding to the block size information is output synchronously during a plurality of clock cycles of the external clock signal."
- 3.8.7. It follows that the subject-matter of claim 2 is not new.
- 3.8.8. As discussed, blocks of data are output from the slave onto the external bus synchronously with respect to CLK* [NuBus, Fig. 5]. The data is output synchronously with respect to the rising edge: "[B]us signals shall be changed only at the rising edge of CLK*" [NuBus, Fig. 5; p. 24 (Section 3.1.8.1)]. Thus, NuBus discloses that "the data is output onto the external bus synchronously with respect to a rising edge transition of the external clock signal."
- 3.8.9. It follows that the subject-matter of claim 8 is not new.
- 3.8.10. In NuBus there are four different block sizes 2, 4, 8, and 16. There are four bits to define the block size (AD2 - AD5). If AD2 = High then the size is 2, AD3 = High then size is 4, etc. [NuBus, Table 3]. Thus, the AD lines represent the block size in binary notation. Thus, NuBus discloses that "the block size information is a binary representation of the amount of data to be output after receipt of a read request."
- 3.8.11. It follows that the subject-matter of claim 10 is not new.
- 3.9. Scalable Coherent Interface ("SCI")
- 3.9.1. "The Scalable Coherent Interface Project (SuperBus)", SCI-22Aug88-doc ("SCI A")
- 3.9.2. "Scalable Coherent Interface", SCI-28Nov88-doc20 ("SCI B")
- 3.9.3. P1596: "SCI, A Scalable Coherent Interface", SCI-28Nov88-doc 2 ("SCI C")
- 3.9.4. "Proposal for Clock Distribution in SCI" - 5/5/89 ("SCI D")
- 3.9.5. Norsk Data Report - "A Proposal for SCI Operation" by Knut Alnes - November 1988 ("SCI E")
- 3.9.6. "Scalable I/O Architecture for Buses" by David V. James, SCI-28Nov88-doc3 ("SCI F")

- 3.9.7. SCI is an interface standard used to facilitate the assembly of nodes, or devices, including processors, I/O devices and memory [SCI A, p. 6]. A node can be a simple memory module [SCI F, Fig. "Board Architecture" p. 3]. SCI uses a 16 bit wide synchronous, packetized bus, to carry address, data and control information [SCI A, p. 4]. Thus, SCI discloses a "semiconductor memory device having at least one memory array which includes a plurality of memory cells" and its method of operation.
- 3.9.8. Each SCI node receives a signal from a central system clock [SCI A, Fig. 3; SCI B, p. 2]. Thus, SCI discloses "receiving an external clock signal having a fixed frequency."
- 3.9.9. An SCI node receives a request packet with target, source, control and address information [SCI A, Figs. 4 and 15]. SCI supports operations of 32, 64, 128 and 256 data blocks, and 1-16 byte subsets of the 16-byte block [SCI A, p. 8; Fig. 11; SCI C, p. 11]. The block size information is conveyed in the transfer code, which forms part of the control information received by the node [SCI A, Figs. 14 and 15; p. 14; SCI C, p. 13 (Header Command)]. Thus, SCI discloses "receiving block size information, wherein the block size information defines an amount of data to be output onto an external bus in response to a read request."
- 3.9.10. The targeted device outputs a variable size data block (i.e., Data Word 0 to Data Word n) in accordance with the block size information [SCI A, Fig. 15; p. 14]. All inputs and outputs on the bus are synchronous with respect to the external system clock [SCI A, p. 3]. All nodes in SCI operate synchronously with respect to the system clock. Thus, SCI discloses "outputting the amount of data corresponding to the block size information, in response to a read request, synchronously with respect to the external clock signal."
- 3.9.11. It follows that the subject-matter of claim 1 is not new.
- 3.9.12. The variable size data block (i.e., Data Word 0 through Data Word n) is output in accordance with the block size information [SCI A, Fig. 15; p. 14]. Successive data words are output during a plurality of clock cycles. All inputs and outputs on the bus are synchronous [SCI A, p. 3]. All nodes in the SCI operate synchronously with respect to the system clock. Thus, SCI discloses that "the amount of data corresponding to the block size information is output synchronously during a plurality of clock cycles of the external clock signal."
- 3.9.13. It follows that the subject-matter of claim 2 is not new.
- 3.9.14. SCI discloses the use of both clock edges for changing data [SCI B, p. 2; SCI D, p. 2]. Thus, SCI discloses that "the data is output onto the external bus synchronously with respect to a rising edge transition of the external clock signal" and that "the data is output onto the external bus synchronously with respect to a falling edge transition of the external clock signal."
- 3.9.15. It follows that the subject-matter of claims 8 and 9 is not new.
- 3.9.16. The "read request" and "block size information" are both contained within a "request packet" [SCI A, pp. 4 and 14; Figs. 9 and 15].
- 3.9.17. It follows that the subject-matter of claim 12 is not new.
- 3.9.18. SCI discloses the use of an internal clock signal generated by a digital phase lock loop to eliminate clock skew [SCI B, p. 2; SCI D, pp. 1-5]. The internal clock is used for outputting data on the external bus [SCI B, p. 2]. Thus, SCI discloses "generating a first internal clock signal using the external clock signal wherein the amount of data corresponding to the block size information is output on to the external bus synchronously with respect to the first internal clock signal."

- 3.9.19. It follows that the subject-matter of claim 13 is not new.
- 3.9.20. As discussed, SCI discloses the use of a digital phase lock loop to generate an internal clock signal [SCI B, p. 2; SCI D, pp. 1-3]. The digital PLL disclosed in the reference D is a DLL; it shows the use of a delay line and no VCO. Thus, "the first internal clock signal is generated by a delay locked loop."
- 3.9.21. It follows that the subject-matter of claim 14 is not new.
- 3.9.22. Each SCI node has a unique 16 bit identification code stored in a register [SCI A, p. 6; SCI E, pp. 1-4]). Thus, SCI discloses that "the semiconductor memory further includes a programmable identification register to store an identification value to identify the memory device from a plurality of other memory devices on the external bus." The ID information is sent in the request packet and received by the SCI nodes [SCI A, pp. 6-7; Fig. 9]. Thus, SCI discloses "receiving identification information."
- 3.9.23. SCI nodes determine whether the target identification code in the request packet matches the identification code of the node [SCI E, pp. 1 and 5-6]. The targeted node outputs a variable size data block (i.e., Data Word 0 though Data Word n) in accordance with the block size information [SCI A, Fig. 15; p. 14]. All inputs and outputs on the bus are synchronous with respect to the external system clock [SCI A, p. 3]. All nodes in the SCI operate synchronously with respect to the system clock. Thus, SCI discloses "determining whether the identification information corresponds to the identification value stored in the programmable identification register wherein, when the identification information corresponds to the identification value, the amount of data corresponding to the block size information is output onto the external bus synchronously with respect to the external clock signal."
- 3.9.24. It follows that the subject-matter of claim 17 is not new.
- 3.9.25. The identification register contains a Global ID and Local ID. The Global ID is the 10 MSB of the node ID and is used to identify a ring on the SCI network (which contains several nodes) [SCI E, p. 1]. Thus, SCI discloses that "the programmable identification register stores an identification value to identify the memory device and a plurality of other memory devices on the external bus."
- 3.9.26. It follows that the subject-matter of claim 18 is not new.
- 3.9.27. Node ID's are assigned during an initialisation sequence after reset of the SCI. A hardware or software based protocol is employed to assign node identifications. Following the assignment of node identification, the SCI master will initialise the interface registers of the various nodes and will cause the node ID to be written into the individual registers [SCI E, pp. 2-4]. Thus, SCI discloses that "the programmable identification register is programmed after power is applied to the memory device duration initialisation of the memory device" and that "the programmable identification register stores an identification value during an initialisation sequence of the memory device."
- 3.9.28. It follows that the subject-matter of claims 19 and 20 is not new.

3.10. US Patent 4,785,394, Published November 15, 1988 ("Fischer")

- 3.10.1. Fischer describes a multiprocessor computer system arranged around a split-transaction bus [Fischer, abstract]. Bus devices are notionally classified into "initiators," including CPU modules, and "responders," including memory modules [Fischer, Fig. 1]. The split-transaction bus carries all information between bus devices and no point-to-point signals are used [Fischer, 7:36-40]. Each of the memory modules includes conventional memory components [Fischer, 6:6-7]. Some

or all of the memories in the modules may be cache memories [Fischer, 6:17-19]. Both the conventional memories and cache memories are semiconductor memories. Each includes semiconductor logic components [Fischer, 5:47-49]. Thus, Fischer discloses a "semiconductor memory device having at least one memory array which includes a plurality of memory cells."

- 3.10.2. A typical responder module is illustrated in Fig. 2 (right hand side) of Fischer. As can be seen, it includes clock receivers 76 that receive clock signals B.CLK 0, B.CLK 1. As shown in Figs. 3A and 3B of Fischer, these clock signals are fixed frequency square waves in quadrature [Fischer, 8:53-55]. They are generated by a backplane clock generation circuit 74. Clock skew is eliminated by the use of equal length clock signal conductors, ensuring reliable synchronisation [Fischer, 8:33-50]. Thus, Fischer discloses "receiving an external clock signal having a fixed frequency."
- 3.10.3. The initiating transaction of a read operation is illustrated in Fig. 6B [Fischer, 10:32-35]. The format of the address information transmitted in bus cycle X of Fig. 6B is shown in Fig. 7A [Fischer, 14:26-29]. It includes three fields. A first field 86 is a two-bit field indicating the nature of the transaction (read, test and set, scrub or write) [Fischer, 13:59-64]. A second field 88 is a 28 bit memory address that identifies the responder to which the request is addressed and the address within that responder [Fischer, 13:64-14:4]. A third field 90 indicates whether one, two or four doublewords are to be transferred in response to the request [Fischer, 14:4-8]. Thus, Fischer discloses "receiving block size information, wherein the block size information defines an amount of data to be output onto an external bus in response to a read request."
- 3.10.4. Each clock receiver 76, illustrated in Fig. 4 of Fischer, generates four internal clock signals B0 ... B3, one quarter cycle out of phase with each other [Fischer, 9:6-25; Figs. 3C-3F]. The format of data placed on the bus in response to a read request is illustrated in Fig. 7B. As shown in Fig. 6B, one doubleword is output per bus cycle. The Fischer system is synchronous [Fischer, 8:33-50; 9:26-38]. Control, address and data signals transferred from bus master to bus slave (initiator to responder or *vice versa*) are designated as signals B.DAT31-0 [Fischer, 15:43-50]. These signals are asserted on the rising edge of B0 and negated on the rising edge of B3 (three quarter-cycles later). B0 is synchronised with external clock signal B.CLK 1 [Fischer, Figs. 3A, 3C and 4]. Thus, Fischer discloses "outputting the amount of data corresponding to the block size information, in response to a read request, synchronously with respect to the external clock signal."
- 3.10.5. It follows that the subject-matter of claim 1 is not new.
- 3.10.6. As discussed above and as shown in Fig. 6B, one doubleword is output per bus cycle. The Fischer system is synchronous [Fischer, 8:33-50; 9:26-38]. Control, address and data signals transferred from bus master to bus slave (initiator to responder or *vice versa*) are designated as signals B.DAT31-0 [Fischer, 15:43-50]. These signals are asserted on the rising edge of B0 and negated on the rising edge of B3 (three quarter-cycles later). B0 is synchronised with external clock signal B.CLK 1 [Fischer, Figs. 3A, 3C and 4]. Thus, Fischer discloses that "the amount of data corresponding to the block size information is output synchronously during a plurality of clock cycles of the external clock signal."
- 3.10.7. It follows that the subject-matter of claim 2 is not new.
- 3.10.8. Again, as discussed above, control, address and data signals B.DAT31-0 are asserted on the rising edge of B0, which is synchronised with external clock signal B.CLK 1 [Fischer, Figs. 3A, 3C and 4]. Thus, Fischer discloses that "data is output onto the external bus synchronously with respect to a rising edge transition of the external clock signal."

- 3.10.9. It follows that the subject-matter of claim 8 is not new.
- 3.10.10. No inventive step would be involved in advancing or retarding all of the clock signals of Fischer by one half clock cycle. Thus, bus cycles would begin and end on the rising edge of B2 rather than B0. In that case, "data is output onto the external bus synchronously with respect to a falling edge transition of the external clock signal."
- 3.10.11. It follows that the subject-matter of claim 9 is obvious.
- 3.10.12. As discussed above, the third field 90 of the address information transmitted in bus cycle X of Fig. 6B indicates whether one [0,0], two [0,1] or four [1,1] doublewords are to be transferred in response to the request [Fischer, 14:4-8]. Thus, the two bits are a logarithmic binary representation of the block size in units of 32 bits (doublewords). Therefore, Fischer discloses that "the block size information is a binary representation of the amount of data to be output after receipt of a read request."
- 3.10.13. It follows that the subject-matter of claim 10 is not new.
- 3.10.14. As discussed above, a two-bit field 90 of the address information transmitted in bus cycle X of Fig. 6B contains the encoded block size information. Thus, Fischer discloses that "the read request and block size information are included in one request packet."
- 3.10.15. It follows that the subject-matter of claim 12 is not new.
- 3.10.16. As discussed above, control, address and data signals B.DAT31-0 are asserted on the rising edge of internal clock signal B0. Thus, Fischer discloses "generating a first internal clock signal using the external clock signal wherein the amount of data corresponding to the block size information is output on to the external bus synchronously with respect to the first internal clock signal."
- 3.10.17. It follows that the subject-matter of claim 13 is not new.
- 3.10.18. As discussed above, the control, address and data signals B.DAT31-0 are negated on the rising edge of internal clock signal B3. Thus, Fischer discloses "generating first and second internal clock signals using the external clock signal wherein the amount of data corresponding to the block size information is output on to the external bus synchronously with respect to the first and second internal clock signals."
- 3.10.19. It follows that the subject-matter of claim 15 is not new.
- 3.10.20. As discussed above, the format of the address information transmitted in bus cycle X of Fig. 6B is shown in Fig. 7A [Fischer, 14:26-29]. It includes three fields. The second field 88 is a 28 bit memory address that identifies the responder to which the request is addressed and the address within that responder [Fischer, 13:64-14:4]. Thus, Fischer discloses "receiving identification information." No point to point signals are used [Fischer, 7:36-40]. Thus, each memory module must decode the MSBs of the address to determine whether to respond. Accordingly, the module must contain an internal address space identification with which to compare the MSBs. That address space identification must be programmable, and in Fischer it is received from the backplane [Fischer, 16:21-27]. Thus, Fischer discloses "a programmable identification register to store an identification value to identify the memory device from a plurality of other memory devices on the external bus" and "determining whether the identification information corresponds to the identification value stored in the programmable identification register wherein, when the identification information corresponds to the identification

value, the amount of data corresponding to the block size information is output onto the external bus synchronously with respect to the external clock signal."

- 3.10.21. It follows that the subject-matter of claim 17 is not new.
- 3.10.22. As discussed above, Fischer describes receiving an address space identification from the backplane. This must be done during an initialisation sequence following power-up or the memory would be unable to respond to requests.
- 3.10.23. It follows that the subject-matter of claims 19 and 20 is not new.

4. LACK OF INVENTIVE STEP ART. 56 EPC

- 4.1. The matter disclosed in the Parent Application does not support any claim purporting to cover buses other than the highly multiplexed bus described or read requests other than read request packets. In proceedings brought in the UK, France and Germany under the patent arising from the Parent Application, the patentee has asserted that the terms "external bus" and "request" in granted claim 1 of that patent should be afforded an interpretation wider than that supported by the Parent Application. The following discussion assumes that such a wider interpretation may be advanced by the patentee in respect of the Patent.
- 4.2. The subject-matter of claims 1-21 of the Patent is obvious even if entitled to the declared priority date of April 18, 1990.

4.3. Other Relevant Documents

4.3.1. Programmable Block Size Reads (Claims 1, 2, 8-10 and 12)

4.3.1.1. Japanese Patent Application No. S63-142445, Published June 14, 1988, and English Translation ("Taguchi")

4.3.1.1.1. Taguchi discloses a memory device having an array formed from a plurality of memory cells [Taguchi, Fig. 1 in conjunction with Figs. 5 and 6 and 4:27-29] and its method of operation.

4.3.1.1.2. A data length register holds the entire size of the data to be accessed [Taguchi, p. 4]. In addition, Taguchi discloses the use of a block length register to define the size of each block and a data length register to define the number of blocks that are to be accessed in an operation [Taguchi, p. 4]. Block data of the size specified in the block length register is read from the memory, the next block of memory to be accessed is calculated from the starting address plus the block size, another block read is performed, and the new starting address is calculated [Taguchi, p. 4]. Counters monitor the number of block reads left in the block access operation. The reading process continues until the amount of data specified in the data length register is output [Taguchi, p. 4].

4.3.1.2. US Patent 4,315,308, Published February 9, 1982 ("Jackson")

4.3.1.2.1. Jackson discloses a system with a microprocessor attached to a Bus Interface Unit (BIU) which provides for interface control of data transfers between the processor and devices, including memory devices [Jackson, 4:17-21]. Each unit in the system receives a clocking signal CLKA to control synchronisation [Jackson 4:36-44; 5:36-38]. CLKA is of fixed frequency [Jackson, Fig. 2].

4.3.1.2.2. Jackson discloses block modes of up to 20 bytes per read request [Jackson, 3:16-19; 6:4-6]. Specifically, Jackson supports transactions of 1, 2, 4, 6, 8, 10, 16 and 20 bytes [Jackson, Fig. 3]. A three bit long

sequence in a control specification is used to specify block size [Jackson, Fig. 3 (bits 10, 11 and 12); 5:29-30]. The memory returns the requested number of data bytes to the BIU [Jackson, 6:8-12]. As the memory continues to output the data, the BIU buffers and aligns the data and then transfers the data across the ACD bus to the processor. This process is repeated until the processor has received the requested number of bytes [Jackson, 6:20-23]. A read operation takes place over a plurality of clock cycles [Jackson, 5:38-39; Fig. 9]. The block size information is a binary representation of the amount of data to be output [Jackson, Fig. 3].

4.3.1.2.3. Jackson discloses the use of a control specification whereby the operation code (read request), block size information and the eight least significant bits of the address are placed on the bus at the same time [Jackson, Figs. 2 and 3; 5:53-58].

4.3.2. Programmable Access Time Register (Claims 3-7 and 11)

4.3.2.1. Japanese Patent Application Sho 62-71428, Published October 5, 1988, and English Translation ("Yamaguchi")

4.3.2.1.1. Yamaguchi describes a dual-port RAM that has both random access I/O and serial access I/O capabilities [Yamaguchi, 3:1-3; 7:3-9]. The random access and serial ports IO1 ... IO3, SIO1 ... SIO 3 and address lines A0 ... A_i are provided to allow the device to connect to an external bus. The RAM includes four memory arrays M-ARY1 ... M-ARY4 [Yamaguchi, 8:10-12]. Each memory cell array M-ARY1 ... M-ARY4 comprises $m+1$ word lines and $n+1$ sets of complementary data lines which intersect at $(m+1) \times (n+1)$ memory cells [Yamaguchi, 8:17-20]. The device may be formed on a single chip [Yamaguchi, 7:4-6].

4.3.2.1.2. The dual-port RAM includes a timing control circuit TC that receives a number of external control signals [Yamaguchi, 18:10-13; Fig. 1]. One such external control signal is an external serial clock signal SC that is generated off-chip [Yamaguchi, 18:13-16]. The external serial clock signal SC is used to ensure stable synchronisation of the serial output operation of the dual-port RAM with the dot rate of a high-resolution, high dot rate external CRT [Yamaguchi, 28:7-15]. The external serial clock signal is therefore a signal of a fixed frequency linked to the dot rate of the external CRT.

4.3.2.1.3. The timing control circuit includes a counter circuit CTR [Yamaguchi, 19:12-15]. A counter is a register that is capable of incrementing or decrementing the value it contains. Many microprocessor internal registers, including program counter (PC) and accumulator (AC) registers are also counters of this kind. An alternative arrangement is discussed in which a register latches the count value and a count-up counter is used to count from zero until its output matches the value in the register [Yamaguchi, 28:19-29:1]. In the preferred embodiment, the counter of Yamaguchi is loaded with a value presented on the parallel I/O lines IO1 ... IO4 on the falling edge of the row address strobe signal RAS_i [Yamaguchi, 21:12-17; 25:8-11]. Once loaded with this value, the counter circuit counts down to zero in synchronism with an internal clock signal, the counter advancing timing signal ϕ_{cp} [Yamaguchi, 21:18-22:3]. This internal clock signal ϕ_{cp} is generated by the timing control circuit TC from, and has the same frequency as, the external serial clock signal SC. It is merely a slightly delayed version of the external serial clock signal SC owing to the gate delays of inverters N1 and N2 and the AND gate AG2 of the timing control circuit TC [Yamaguchi, 23:8-14, Figs. 1 & 3]. The overall result is that the counter

circuit CTR counts clock cycles of the external serial clock signal SC [Yamaguchi, 19:12-15].

4.3.2.1.4. As will be described below, the counter circuit is used to delay the output of serial data by a number of clock cycles corresponding to the value loaded into the counter circuit CTR [Yamaguchi, 27:19-21].

4.3.2.1.5. In accordance with a further internal clock signal ϕ_c , the serial I/O circuit latches the data presented by data registers DR1 ... DR4 on complementary data lines CDS1 ... CDS4 to the serial data lines SIO1 ... SIO4 and hence to an external bus [Yamaguchi, 19:17-19]. Clock signal ϕ_c is synchronised to external serial clock signal SC [Yamaguchi, 26:17-19]. After the first transition of the internal clock signal ϕ_c , it is also used to shift a shift register of a pointer PNT that points to the current position in the data registers DR1 ... DR4, thus accessing the next data in the data registers DR1 ... DR4 [Yamaguchi, 26:19-27:7]. Data is output in accurate synchrony with the external serial clock signal SC only once an internal strobe signal ϕ_{dt} , generated when the counter CTR reaches to zero, has been asserted.

4.3.2.1.6. The serial I/O circuit drives data presented on the complementary data lines CDS1 ... CDS4 to the serial data lines SIO1 ... SIO4 in accordance with the internal clock signal ϕ_c . One or other edge must be used to control SIO. Fig. 3 shows that the output data transitions on the rising edge of ϕ_c and on the rising edge of the serial clock signal SC.

4.3.2.1.7. As explained above, the number of clock cycle delays introduced into a serial read operation by the counter circuit CTR is equal to the value stored in it during the read request. That value can be between 0 and 15. The value stored is therefore representative of one of a plurality of different delay times.

4.3.2.2. GigaBit Logic, 1989 GaAs IC Data Book & Designer's Guide, August 1989, 12G014 256x4-Bit Registered Self-Timed SRAM ("GigaBit")

4.3.2.2.1. Gigabit describes a 256x4-bit static RAM fabricated using Gallium Arsenide (GaAs) technology ("The 12G014"). The 12G014 is a self-timed SRAM ("STRAM"), running at a fixed frequency of 400 MHz. The clock cycle is therefore 2.5 ns. The 12G014 SRAM has differential clock inputs CLK and CLK₁, as shown in the block diagram on page 2-3. The differential clock inputs connect to an output clock generator functional block, which must include clock receiver circuitry to receive the differential clock inputs.

4.3.2.2.2. The 12G014 is a $\times 4$ memory, with four output drivers (Q0-Q3) outputting data onto the bus in response to a read request. The output is fully registered (double latch) as shown in the block diagram on page 2-3. The output register is clocked synchronously to the external clock via the agency of the output clock generator that generates an internal clock from the complementary clock inputs CLK and CLK₁. Hence, the output operation is done synchronously with respect to the external clock.

4.3.2.2.3. The 12G014 has three output modes, namely latch mode, register mode and transparent mode. The transparent mode is asynchronous (similar to conventional SRAMs). Both the latched mode and register mode are clocked. In register mode, memory access takes place during the clock cycle in which the read request is received. Valid output data is presented to the on-chip output drivers. At the next rising edge of the

clock signal, i.e. a full 2.5 ns clock cycle after the read request is received, this data is loaded into the output drivers and propagates to the data output pins, and hence the bus, where it is held for a full cycle. This is shown and described in the timing diagram on page 2-6.

- 4.3.2.2.4. In latch mode, memory access again takes place during the clock cycle in which the read request is received. Valid output data is presented to the on-chip output drivers. However, these output drivers are driven transparent at the falling edge of the clock signal to allow valid data to appear on the data output pins, and hence on the bus, as soon as possible. The duty cycle of the clock is changed so that the falling edge arrives after less than one half clock cycle. The output drivers are driven transparent at this falling edge and latched at the next rising edge to hold data over to the falling edge of the next clock cycle. This is shown and described in the timing diagram on page 2-7.
- 4.3.2.2.5. The output mode is programmed by applying one of three signal levels to a MODE pin, as described on page 2-4. Vss level gives register mode, Vdd gives latch mode and Vcc gives transparent mode. The mode pin signal is received in the output clock generator where it must be decoded. The circuitry that decodes the mode pin signal and produces the decoded logical outputs is programmable until the mode pin is connected and thereafter outputs the value programmed into it. The value programmed into this register determines the output mode and, in particular, the output delay.
- 4.3.2.2.6. As discussed above, in register mode, data output takes place on the rising edge of the internal clock, which is synchronised with the external clock signal CLK. The Fig. on page 2-3 includes an "Output Clock Generator," which generates internal clock signals from CLK and CLK¹. These internal clock signals drive the output register. In register mode, data is driven onto the bus at the rising edge of the internal clock signal.
- 4.3.2.2.7. As discussed above, the value programmed into the access-time register determines the output mode and, in particular, the output delay. It is representative of one of a plurality of different delay times.

4.3.2.3. US Patent No. 4,499,536, Issued February 12, 1985 ("Gemma")

- 4.3.2.3.1. Gemma describes a processor-based SCU that interfaces with a main memory including a plurality of memory cells arranged as one or more memory arrays [Gemma, 3:65]. The memory arrays may be in a number of configurations [Gemma, 6:25-32]. The SCU includes a controller 17 that receives a clock signal T₂ synchronised with the processor machine cycle [Gemma, 4:26-35; Fig. 2]. This is received by AND gates 29 and 24 within the controller 17. Timing control signals that drive a counter 25 are generated from the clock signal T₂ [Gemma, Fig. 2].
- 4.3.2.3.2. The SCU 17 receives memory related instructions from a current instruction register via a signal line 101 under the control of an execution unit. Control information is sent to the SCU from the execution unit via a signal line 102. The control information determines the "type of main memory access, that is, read (FE), full write (ST) and partial write (PST)" [Gemma, 3:10-17]. The generation of a main memory access start signal (EX) that is sent to the main memory is described. The signal has a pulse width predetermined by the timing of timing signals T₀ and T₁ [Gemma, 3:18-41]. Subsequently, the SCU 19 sends a GFDR signal to the main memory on signal line 19 to instruct the main memory to send readout data to a data bus 118 and sends an

ADV signal to the processor on signal line 120 to instruct a read data register in the processor to read the data on the data bus 118 [Gemma, 3:52-56]. The content of the read data register is sent to the execution unit 2 via a data bus 122 [Gemma, 3:58-60]. It is implicit in the requirement for the main memory to respond to the GFDR signal that the main memory includes output drivers that respond to the signal. These output drivers output data onto the data bus 118.

4.3.2.3.3. The SCU includes a controller 13, which is shown in detail in Fig. 2. The controller 13 has "configuration registers 20 and 21 which retain identification flags for the machine cycle of the processor and identification flags concerning the access time of the main memory cells, respectively" [Gemma, 3:61-65]. As described, the controller includes "a counter 25 for counting an elapsed time after the signal EX of a predetermined pulse width has been produced, a decoder 26 for decoding the count of the counter 25, a control circuit 27 for producing the signals GFDR and ADV ... based on the output of the decoder 26" [Gemma, 4:4-10]. Therefore, the timing of the GFDR and ADV signals is determined with reference to the start of the count of the counter. The count is begun with reference to the EX signal that has a predetermined pulse width [Gemma, 3:39-41]. The counter counts according to the clock signal T_2 synchronised with the machine cycles [Gemma, 4:26-35].

4.3.2.3.4. The configuration registers 20 and 21 of the controller 13 are initialised to contain one of three machine cycle identification flags and one of three main memory identification flags [Gemma, 4:11-17]. The output from the configuration registers 20 and 21 are supplied to the control circuit 22 via signal lines 200 and 201 [Gemma, 4:22-25]. The counter 25 is reset in response to the memory signal EX being low [Gemma, 4:32-33]. The output of the counter, representing a number in binary form, 20-23, is fed to a decoder 26 [Gemma, 4:35-38].

4.3.2.3.5. According to Gemma, "the relation between the combinations of the identification flags m_k ($k=1-3$) and the t_j ($j=1-3$) and the send timings of the signals GFDR, ADV and BSYR. The send timing are represented by the counts C_i-C_{i+6} of the number of machine cycle counted after the signal EX has been produced. Those counts are predetermined based on the response performance of the processor and the main memory, i.e. based on the contents of the configuration registers [Gemma, 4:53-60]. The following example of the timing is given: "when the identification flags m_1 and t_1 are set to "1", respectively, the signal GFDR is sent at the timing C_i and the signal ADV is sent at the timing C_{i+1} " [Gemma, 4:60-63]. Thus, once the EX signal has been produced, that is, once the memory has been instructed to execute a read, the signal GFDR is sent to the memory to instruct it to output the result of the read at a time of C_i which is measured in terms of machine cycles of the processor, the number being determined by the contents of the configuration registers. Having issued the signal GFDR, the controller then issues the signal ADV at a time of C_{i+1} to instruct the read data register to read the output from the memory.

4.3.2.3.6. As discussed above, the timing signal T_2 is synchronous with the machine cycles. When signal EX goes high, AND gate 24 detects the rising edge of signal T_2 . Thus, counter 25, and decoder 26 change state on the rising edge of T_2 . C_i transitions from low to high [Gemma, 4:41-44]. The same will happen with $C_{i+1}-C_{i+6}$. Data selecting circuits 22 select, according to the input $s0-s4$, one of the signals C_i-C_{i+6} applied to input terminals $d0-d4$. The data selecting circuits reproduce the selected signal at a corresponding output terminal [Gemma 5:48-53].

Unless the data selecting circuits are enabled, the outputs are zero [Gemma 5:60-62].

- 4.3.2.3.7. Thus, the output "u" of a data selecting circuit reflects the state of a corresponding selected input, $C_{i+1} \sim C_{i+6}$. Since the signals $C_{i+1} \sim C_{i+6}$ are positive going and the disabled output of the selecting circuits is zero, the outputs "u" are also positive going in response to a rising edge of T_2 , which is in turn synchronised to an external clock. The signal GFDR, which controls the output of the data from the read data register, is derived from the second data selecting circuit 66.
- 4.3.2.3.8. The initialisation of the configuration registers 20 and 21 is described as follows. In the initialisation of the processor, one of the identification flags $m_1 \sim m_3$ and one of the identification flags $t_1 \sim t_3$ are set to "1". The operation may be carried out by a known technique, such as by loading a microprogram into the control memory during the initialisation of the processor [Gemma, 4:9-22].
- 4.3.2.3.9. The values, $m_1 \sim m_3$ and $t_1 \sim t_3$, stored in the configuration control registers clearly control the timing of the output signals. The value stored in the programmable access time register is representative of one of a plurality of different delay times.
- 4.3.2.4. Japanese Patent Application Sho 62-185253, Published January 31, 1989, and English Translation ("Kumagai")
- 4.3.2.4.1. Kumagai discloses a main storage unit MS 4 that includes RAM memory arrays RAM0, RAM1, RAM2, RAM3 [Kumagai, Fig. 1]. The memory device MS 4 is clocked by an external clock source 5. That clock source is at least common to the memory device MS 4 and the memory controller SCU. It is a fixed frequency clock [Kumagai, Fig. 4]. The clock is used to clock all the interfaces of MS 4 which are all latched: command/address buffer MRQ 20, input data MSD 21 and output data MFD 30 [Kumagai, 4:8-22]. The circuitry to receive the external clock is not specified in Kumagai, but it must exist and can simply be an input buffer.
- 4.3.2.4.2. The control unit of memory device MS 4, MCR 51, contains the clock counter circuitry shown in Fig. 3. The circuitry includes latches C0 ... C3. These latches store a value representative of a number of clock cycles of the external clock, which is input to the circuitry at 316 (clock T0). The combination of blocks 300, 301, 318 and 319 of Fig. 3 is a clock counter that would signal (see input to block 102) when the clock count reaches the delay value stored in latches C0-C3. That signal then is used to trigger CAS1 (through the clock phase selection circuitry 102, 303, 304 and 305). Hence, data in RAM0-3 is accessed only after the number of clock cycles stored in C0-C3 has transpired. The data is then sent to the memory output latch MFD 30 to be output at a clock edge to the requesting device SCU 3 [Kumagai, 5:10-27]. Latches C0-C3 are programmable in the sense that they hold whatever values were read into them. Because C0-C3 are part of MCR 51, which is programmed by SCU 3 via MRQ 20, C0-C3 are programmable by the SCU 3 via that interface.
- 4.3.2.4.3. Each of the synchronous interfaces to the memory device MS 4 of Kumagai, namely the command/address interface at MRQ 20, the input data interface at MSD 21 and the output data interface MFD 30 can each be one or many bits wide. That output operation is in response to a read request from SCU 3 via command interface MRQ 20 [Kumagai, 4:32-5:2]. The output operation is delayed until after a specified number

of clock cycles has transpired. The output operation is synchronous with respect to the external clock because MFD 30 is a latch.

4.3.2.4.4. As shown in Fig. 3 and 4 of Kumagai, the memory device MS 4 can be programmed to output data at any one of four phases of the external clock. The phase number value is stored in latches T0-T3 (310-313). One of those four values (e.g. 3 out of a 0-3 range) corresponds to a synchronous output that is synchronised with respect to a rising edge of the external clock signal.

4.3.2.4.5. The value for access time in Kumagai is stored in latches C0-C3. Latches are devices that hold values clocked in after power is applied. The latches C0-C3 are programmed via the command interface MRQ 20. In most cases, it would be done once and for all after power up, because in a given system the memory access time and processor machine cycle will not change. This amounts to programming during an initialisation sequence of the memory device following power up. The latches C0-C3 of Kumagai allow four different access times to be programmed into the memory device MS 4 [Kumagai, 5:27, Fig. 3: blocks 300, 318, 319 and 314].

4.3.3. Delay Locked Loop (claims 14 and 16)

4.3.3.1. At the priority date of the Patent, DLLs were common general knowledge. DLLs enable regularly cyclic digital signals at the same frequency to be synchronised with one another.

4.3.3.2. UK Patent Application GB-2,197,553, Published May 18, 1988 ("Lofgren")

4.3.3.2.1. Lofgren describes a digital phase locked loop circuit (DLL) [Lofgren, abstract]. One principal application of the DLL described is to provide "optimum timing for control of high speed dynamic RAM devices" [Lofgren, 1:14-18]. Lofgren discloses the use of two identical delay lines. One delay line 12 is clocked by a local oscillator 20 and used to calibrate the DLL. The other delay line 18 is used to provide an accurate delay to an input signal [Lofgren, 1:130-2:5].

4.3.3.2.2. Each delay line 12, 18 consists of plural delay elements D_n , the amount of delay introduced by each of which being controlled by the level of two control signals VCP, VCN, generated by a charge pump 16 in response to phase errors detected in the first delay line 12 [Lofgren, 3:78-114; 4:23-33; 4:62-68; 5:78-96]. The delay introduced by the first delay line 12 is varied until it is exactly one period of the local oscillator 20 [Lofgren, 4:104-106]. Identical control signals are then applied to the second delay line 18. It is then known that each of the n delay elements D_n in the second delay line 18 will introduce a delay of one n^{th} of the period of the local oscillator. The second delay line is a multi-tap line, allowing one of a plurality of delays to be selected [Lofgren, 2:5-13].

4.3.3.3. IEEE Journal of Solid State Circuits, Vol. 25, No. 1, February 1990, "An On-Chip Smart Memory for a Data-Flow CPU" ("Uvieghara"), as Exemplifying Common General Knowledge

4.3.3.3.1. Uvieghara describes a high performance substrate CPU that has an embedded smart memory of the type known as a "register alias table" ("RAT"). The RAT is a multi-port content-addressable memory supporting branch prediction and exception handling. An experimental 1240 bit RAT is described [Uvieghara, abstract]. The RAT is a

synchronous DRAM [Uvieghara, Fig. 5] having a block read mode [Uvieghara, 87:2-4].

4.3.3.3.2. The RAT uses a PLL-based clock generator on-chip [Uvieghara, Fig. 13; and p. 92 "All clocks are generated by a ... PLL-based clock generator."]

4.3.3.4. US Patent No. 4,637,018, Issued January 13, 1987, ("Flora"), as Exemplifying Common General Knowledge

4.3.3.4.1. Flora discloses a DLL that is designed to synchronise the outputs of a distributed clock driver circuit with an accurately delayed external clock signal [Flora, 3:6-16]. As the device is a clock distribution system, exact synchronisation with the external clock signal is not essential, but synchronisation of the outputs of several chips with each other is [Flora, 3:22-28]. For this reason the external clock signal receiving line is bifurcated. One branch is subject to an accurate delay and the other includes a multi-tap delay line introducing a variable delay to the clock driver circuit. The outputs of the clock driver circuitry are synchronised to the delayed external clock [Flora, 3:60-68]. However, if accurate phase synchronisation with the external clock were required, the delay line would be removed, or re-introduced in the same branch as the variable delay to give an accurate delay of just short of one clock cycle. In this way, the outputs of the clock driver circuitry would be accurately synchronised to the external clock. This all results from the application of common general knowledge at the priority date.

4.3.3.5. Japanese Patent Application JP-A-01-284132, published November 15, 1989, and English Translation ("Kosugi"), as Exemplifying Common General Knowledge

4.3.3.5.1. Kosugi discloses a digital phase locked loop (DPLL) that is used for synchronising the internal read clock 3 (clocking the output) of a memory device 1 to its internal write clock 2 (clocking the input). The memory write clock 2 is itself generated from an external clock.

4.3.3.6. Motorola MC88200 Cache/Memory Management Unit User's Manual, Published 1988, as Exemplifying Common General Knowledge ("MC88200")

4.3.3.6.1. The MC88200 is a single chip synchronous device that contains high-speed cache memory. The MC88200 has a large number of on-chip, dynamically programmable configuration registers, as shown in table 1-1. The device generates on-chip all internal timing signals from an external clock signal CLK. The MC88200 internal clock is normally phase locked to the external clock signal CLK in order to minimise the skew between the external and internal signals [MC88200, page 4-9].

4.3.3.7. It was obviously desirable on the priority date of the Patent that internal device clock signals should be synchronised as closely as possible with the external clock signal.

4.3.3.8. The objective problem to be solved by a device according to granted claims 14 or 16 is the provision of more accurately synchronised internal clock signals. This problem is solved by using a DLL, as would have been well known to a person skilled in the art at the priority date.

4.3.3.9. It follows that the subject-matter of claims 14 and 16 is obvious.

4.3.4. Dual Internal Clocks (Claim 15)**4.3.4.1. US Patent 4,680,738, Published July 14, 1987 ("Tam")**

4.3.4.1.1. Tam discloses a memory device that can operate in a conventional access mode and a high speed sequential (block) mode [Tam, abstract]. To facilitate the high speed sequential mode, the memory is organised into two arrays [Tam, Fig. 1]. A separate internal clock is derived for each array (CKL; CKR) from an external clock signal CK. The two internal clocks are 180° out of phase with one another (Tam, Fig. 2; 3:63-4:5 *et seq.*).

4.3.4.2. US Patent 4,330,852, Published May 18, 1982 ("Redwine")

4.3.4.2.1. Redwine discloses a serial input/output memory device in which the input/output is performed in units of 256 bits [Redwine, abstract; Fig. 1]. To increase the serial access speed, the memory is organised into two arrays [Redwine, 2:55-60; Fig. 1]. A separate internal clock is derived for each array (Φ_1 ; Φ_2) from an external clock signal Φ . The two internal clocks are 180° out of phase with one another [Redwine, 4:62-64; Fig. 3]. Thus, the memory device is able to output serial data via output multiplexer 26 at twice the clock rate of the external clock Φ [Redwine, 7:17-26].

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
(Case No. P043D2C3C)

In the Application of:)
FARMWALD et al.)
Serial No: 09/492,982) Group
Filed: JANUARY 27, 2000) Art Unit:2818
Title: METHOD OF OPERATING A MEMORY)
DEVICE HAVING A VARIABLE DATA)
INPUT LENGTH) Before
Examiner: T. Nguyen

Box CPA

FACSIMILE COVER SHEET

DATE: January 30, 2001 FAX RECEIVED
TO: Office of Petitions JAN 30 2001
FAX NUMBER: 703-308-6916 PETITIONS OFFICE
NO. OF PAGES: 10 attached.

MESSAGE

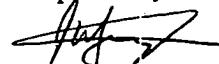
Attached is a petition under 37 CFR 1.313(b)(5), along with a 37 CFR 1.53(d) (CPA) application,

I hereby certify that the attached:

1. Transmittal-Request for a Continued Prosecution Application (CPA)
(2 pages in duplicate)
2. Petition under 37 CFR 1.313(b)(5) For Withdrawal From Issue So That Information
Can Be Considered in an Information Disclosure Statement
(2 pages in duplicate + 2 page attachment)

are being facsimile transmitted to the United States Patent and Trademark Office (Fax No.703-308-6916)
on January 30, 2001 in the above-referenced application.

Respectfully submitted,



Joe G. Moniz
650-947-5336

PRINT OF DRAWINGS
AS ORIGINALLY FILED

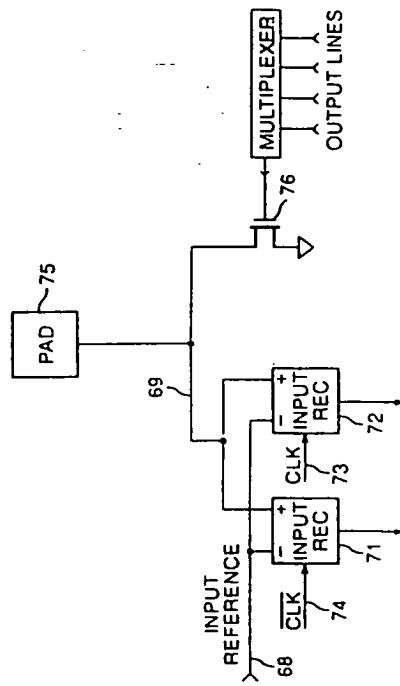


FIG. 10